arXiv:1601.07947v1 [stat.ML] 28 Jan 2016

Large-scale Kernel-based Feature Extraction via Budgeted Nonlinear Subspace Tracking

Fatemeh Sheikholeslami, *Student Member, IEEE*, Dimitris Berberidis, *Student Member, IEEE*, and Georgios B. Giannakis, *Fellow, IEEE*

Abstract—Kernel-based methods enjoy powerful generalization capabilities in handling a variety of learning tasks. When such methods are provided with sufficient training data, broadly-applicable classes of nonlinear functions can be approximated with desired accuracy. Nevertheless, inherent to the *nonparametric* nature of kernel-based estimators are computational and memory requirements that become prohibitive with large-scale datasets. In response to this formidable challenge, the present work puts forward a *low-rank, kernel-based, feature extraction* approach that is particularly tailored for online operation, where data streams need not be stored in memory. A novel generative model is introduced to approximate high-dimensional (possibly infinite) features via a low-rank nonlinear subspace, the learning of which leads to a direct kernel function approximation. Offline and online solvers are developed for the subspace learning task, along with affordable versions, in which the number of stored data vectors is confined to a predefined budget. Analytical results provide performance bounds on how well the kernel matrix as well as kernel-based classification and regression tasks can be approximated by leveraging budgeted online subspace learning and feature extraction schemes. Tests on synthetic and real datasets demonstrate and benchmark the efficiency of the proposed method when linear classification and regression is applied to the extracted features.

Index Terms—Online feature extraction, kernel methods, classification, regression, budgeted learning, subspace tracking.

1 Introduction

KERNEL-BASED expansions can boost the generalization capability of learning tasks by powerfully modeling nonlinear functions, when linear functions fall short in practice. In fact, the upshot of kernel machines is that when provided with sufficient training data, arbitrary nonlinear functions can be approximated with desired accuracy. Although "data deluge" sets the stage by providing the "data-hungry" kernel methods with huge datasets, limited memory and computational capabilities prevent such tools from fully exploiting their learning capabilities. In particular, given a dataset with N vectors of dimension D, a standard kernel regression or classification tool takes $\mathcal{O}(N^2D)$ operations to form the $N \times N$ kernel matrix $\mathbf{K}$, memory $\mathcal{O}(N^2)$ to store it, and $\mathcal{O}(N^3)$ operations to find the optimal pattern.

In this context, several efforts have been made in different fields of stochastic and numerical optimization, functional analysis, and numerical algebra to speed up kernel machines for "big data" applications [8], [11], [18], [22], [27], [32], [38]. A common approach for scaling up kernel methods is to approximate the kernel matrix $\mathbf{K}$ by a low-rank factorization; that is $\mathbf{K} \simeq \mathbf{Z}^\top\mathbf{Z}$, where $\mathbf{Z} \in \mathbb{R}^{r\times N}$ with r $(\ll N)$ is the reduced rank, through which storage and computational requirements go down to $\mathcal{O}(Nr)$ and $\mathcal{O}(Nr^2)$, respectively. Such a factorization can be effected by randomly selecting r training vectors to approximate the kernel matrix as $\mathbf{K} \simeq \mathbf{K}_{\mathcal{S}}^\top\mathbf{K}_{\mathcal{S}}$, where $\mathbf{K}_{\mathcal{S}} \in \mathbb{R}^{N\times r}$ is formed by a subset $\mathcal{S}$ of columns of $\mathbf{K}$ corresponding to the selected data vectors [20]. Nystrom approximation further improves this approach by approximating the kernel matrix as $\mathbf{K} \simeq \mathbf{K}_{\mathcal{S}}^\top\mathbf{W}^\dagger\mathbf{K}_{\mathcal{S}}$, where $\mathbf{W}$ is an $r \times r$ submatrix of $\mathbf{K}$ corresponding to the kernel values computed between the selected r data vectors [38]. In [11], [19], [40], further attempts to improve the Nystrom approximation are made for ensemble datasets, while [1], [7], [41] analyze how such low-rank approximations influence the accuracy of the target tasks.

Feature approximation is a second possibility where instead of approximating the kernel matrix, one directly approximates through an $r \times 1$ feature vector $\mathbf{z}$ the high-dimensional features given by the mapping $\boldsymbol{\phi} : \mathbf{x} \to \boldsymbol{\phi}(\mathbf{x})$, from which the kernel is induced as $\kappa(\mathbf{x}_i, \mathbf{x}_j) =< \boldsymbol{\phi}(\mathbf{x}_i), \boldsymbol{\phi}(\mathbf{x}_j) >$ [21], [22], [27], [39]. Based on $\mathbf{z}$, the nonlinear kernel $\kappa(\mathbf{x}_i, \mathbf{x}_j)$ is approximated by the linear one, meaning $\kappa(\mathbf{x}_i, \mathbf{x}_j) \simeq \mathbf{z}_i^\top\mathbf{z}_j$. Exploiting the fast linear learning machines [12], [32], the kernel-based task then reduces to learning a linear function, which can be achieved in $\mathcal{O}(Nr)$ operations. A third approach to alleviate the prohibitive requirements of kernel methods is offered by online algorithms. Instead of loading the entire datasets in memory, online methods iteratively pass over the set from an external memory [4], [16], [17], [18], [31], [32], [33]. This is also critical when the entire dataset is not available beforehand, but is acquired one datum at a time. For large data streams however, as the number of data increases with time, the number of support vectors (SV) through which the function is estimated, i.e., the set $\mathcal{S}$ in the approximation $f(\mathbf{x}) \simeq \hat{f}(\mathbf{x}) = \sum_{i\in\mathcal{S}} \alpha_i\kappa(\mathbf{x}_i, \mathbf{x})$, also increases. Thus, the function evaluation delay as well as the required memory for storing the SV set eventually become unaffordable. Efforts have been devoted to reducing the number of SVs while trying to

• *This work was supported by NSF grants 1343248, 1500713, and the NIH grant no. 1R01GM104975-01. Preliminary parts of this work were presented at the Proc. of Globalsip Conf., Orlando, FL, Dec. 2015.*
Authors are with the Dept. of Electrical and Comp. Engr. and the Digital Tech. Center, University of Minnesota, Minneapolis, MN 55455, USA. E-mails: {sheik081,bermp001,georgios}@umn.edu

maintain performance on the unseen data (a.k.a. generalization capability) [10]. In more recent attempts, by restricting the maximum number of SVs to a predefined *budget* B, the growth of algorithm complexity is restricted to an affordable limit, and is maintained throughout the online classification [9], [35], [36] or regression [34] task.

The present work introduces a generative model according to which the high-dimensional (possibly infinite) features are approximated by their projection onto a low-rank nonlinear subspace, thus providing a direct kernel function approximation (Section 2). In contrast to [11], [21], [27], [40] where, due to the nature of randomization, the number of required features for providing an accurate kernel function approximation is often large, systematically learning the ambient nonlinear subspace can provide an accurate approximation through a smaller number of extracted features. Offline and online solvers for the subspace learning task are developed in Sections 3 and 4 respectively, and their convergence is analyzed. Furthermore, to keep the complexity and memory requirements affordable, *budgeted* versions of the proposed algorithms are devised in Section 5, in which the number of stored data vectors is confined to a predefined budget B. Budget maintenance is performed through a greedy approach, whose effectiveness is corroborated through simulated tests. Analytical results in Section 6 provide performance bounds on how well the kernel matrix as well as kernelized classification and regression tasks can be approximated by leveraging the novel budgeted online subspace-learning and feature-extraction approach. Finally, Section 7 presents classification and regression experiments on synthetic and real datasets, demonstrating the efficiency of the proposed methods in terms of accuracy and run time.

2 Preliminaries and Problem Statement

Consider a set of N real data vectors of size $D \times 1$. As large values of D and N hinder storage and processing of such datasets, extracting informative features from the data (a.k.a. *dimensionality reduction*) results in huge savings on memory and computational requirements. This fundamentally builds on the premise that the informative part of the data is of low dimension $r < D$, and also the data $\{\mathbf{x}_\nu\}_{\nu=1}^N$ are well represented by the generative model

$$\mathbf{x}_\nu = \mathbf{L}\mathbf{q}_\nu + \mathbf{v}_\nu \ , \quad \nu = 1, \ldots, n \tag{1}$$

where the tall $D \times r$ matrix $\mathbf{L}$ has rank $r < D$; vector $\mathbf{q}_\nu$ is the $r \times 1$ projection of $\mathbf{x}_\nu$ onto the column space of $\mathbf{L}$; and $\mathbf{v}_\nu$ denotes zero-mean noise.

Pursuit of subspace $\mathbf{L}$ and the low-dimensional features $\{\mathbf{q}_\nu\}_{\nu=1}^N$ is possible using a blind least-squares (LS) criterion regularized by a rank-promoting term using e.g., the nuclear norm of $\mathbf{L}\mathbf{Q}_N$, where $\mathbf{Q}_N := [\mathbf{q}_1, ..., \mathbf{q}_N]$ [28]. To arrive at an efficient subspace tracker, the nuclear norm can be further surrogated by the Frobenious-norms of $\mathbf{L}$ and $\mathbf{Q}_N$, yielding

$$\min_{\mathbf{L}, \{\mathbf{q}_\nu\}_{\nu=1}^N} \frac{1}{2N} \sum_{\nu=1}^{n} \|\mathbf{x}_\nu - \mathbf{L}\mathbf{q}_\nu\|_2^2 + \frac{\lambda}{2N} \Big(\|\mathbf{L}\|_F^2 + \|\mathbf{Q}_N\|_F^2 \Big) \tag{2}$$

where λ controls the tradeoff between LS fit and rank regularization [24]. Principal component analysis (PCA) - the "workhorse" of dimensionality reduction- solves (2) when the rank regularization is replaced with orthonormality constraints on the columns of $\mathbf{L}$. Undoubtedly, the accuracy of any linear dimensionality reduction method is dictated by how well the model (1) fits a given dataset, which is related to how well the corresponding data covariance matrix can be approximated by a low-rank matrix [15, p. 534].

In practice however, low-rank linear models often fail to accurately capture the underlying characteristics of the datasets. A means to deal with nonlinearities in pattern recognition tasks, is to first map vectors $\{\mathbf{x}_\nu\}_{\nu=1}^N$ to a higher $\bar{D}$-dimensional space using a function $\boldsymbol{\phi} : \mathbb{R}^D \to \mathbb{R}^{\bar{D}}$ (possibly with $\bar{D} = \infty$), and subsequently seek a linear function is sought over the lifted data $\boldsymbol{\phi}(\mathbf{x})$. Corresponding to the mapping, a so-termed kernel function $\kappa(\mathbf{x}_i, \mathbf{x}_j) = \boldsymbol{\phi}^\top(\mathbf{x}_i)\boldsymbol{\phi}(\mathbf{x}_j)$ is induced, which chosen to have a closed form expression, circumvents the need to explicitly know $\{\boldsymbol{\phi}(\mathbf{x}_\nu)\}_{\nu=1}^N$ - what is referred to as the "kernel trick". Upon defining the $\bar{D} \times N$ matrix $\boldsymbol{\Phi}_N := [\boldsymbol{\phi}(\mathbf{x}_1), ..., \boldsymbol{\phi}(\mathbf{x}_N)]$, the $N \times N$ kernel matrix related to the covariance of the lifted data, is expressed with (i, j) entry $\kappa(\mathbf{x}_i, \mathbf{x}_j)$ as the $\mathbf{K}(\mathbf{x}_{1:N}, \mathbf{x}_{1:N}) = \boldsymbol{\Phi}_N^\top \boldsymbol{\Phi}_N$, where $\mathbf{x}_{1:N} := [\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_N]$. Its computation and storage are $\mathcal{O}(N^2 D)$ and $\mathcal{O}(N^2)$ respectively, which is often not affordable in large datasets ($N \gg$).

In practical large datasets however, $\mathbf{K}$ is approximately low rank. This fact is exploited in e.g., [11], [39] and [38] to approximate $\mathbf{K}$ via a low-rank factorization, hence alleviating the evaluation and memory requirements of offline kernel-based learning tasks from $\mathcal{O}(N^2)$ down to $\mathcal{O}(Nr)$. Here, we further build on this observation to deduce that the low-rank property of $\mathbf{K} = \boldsymbol{\Phi}_N^\top \boldsymbol{\Phi}_N$ implies that $\boldsymbol{\Phi}_N$ can also be approximated by a low-rank matrix, thus motivating our pursuit of *online low-rank factorization* of $\boldsymbol{\Phi}_N$. To this end, instead of projecting $\{\mathbf{x}\}$s onto the columns of $\mathbf{L}$ as in (2), we will project $\{\boldsymbol{\phi}(\mathbf{x})\}$s on $\bar{\mathbf{L}} \in \mathbb{R}^{\bar{D} \times r}$ whose columns span what we refer to as "virtual" column subspace since $\bar{D}$ can be infinite. Specifically, we consider [cf. (2)]

$$\min_{\bar{\mathbf{L}}, \{\mathbf{q}_\nu\}_{\nu=1}^N} \frac{1}{2N} \sum_{\nu=1}^{N} \|\boldsymbol{\phi}(\mathbf{x}_\nu) - \bar{\mathbf{L}}\mathbf{q}_\nu\|_2^2 + \frac{\lambda}{2N} \Big(\|\bar{\mathbf{L}}\|_F^2 + \|\mathbf{Q}_N\|_F^2 \Big) \tag{3}$$

where $\|\bar{\mathbf{L}}\|_F := \sqrt{\sum_{c=1}^r \|\mathbf{l}_c\|_{\mathcal{H}}^2}$ with $\|.\|_{\mathcal{H}}$ denoting the norm defined in the reproducing kernel Hilbert space (RKHS) induced by the kernel $\kappa(\mathbf{x}, .)$, and $\mathbf{l}_c$ denoting the c-th column of $\bar{\mathbf{L}}$. For a fixed $\mathbf{Q}_N$, the criterion in (3) is minimized by the subspace

$$\bar{\mathbf{L}}_N = \boldsymbol{\Phi}_N \mathbf{Q}_N^\top \Big(\mathbf{Q}_N \mathbf{Q}_N^\top + \lambda \mathbf{I} \Big)^{-1} := \boldsymbol{\Phi}_N \mathbf{A} \tag{4}$$

where the $N \times r$ factor $\mathbf{A}$ can be viewed as "morphing' the columns of $\boldsymbol{\Phi}_N$ to offer a flexible basis for the lifted data. Substituting (4) back into (3) and exploiting the kernel trick, we arrive at

$$\begin{aligned} \min_{\mathbf{A}, \{\mathbf{q}_\nu\}_{\nu=1}^N} \quad & \frac{1}{2N} \sum_{\nu=1}^{N} \Big(\kappa(\mathbf{x}_\nu, \mathbf{x}_\nu) - 2\mathbf{k}^\top(\mathbf{x}_{1:N}, \mathbf{x}_\nu)\mathbf{A}\mathbf{q}_\nu \\ & + \mathbf{q}_\nu^\top \mathbf{A}^\top \mathbf{K}(\mathbf{x}_{1:N}, \mathbf{x}_{1:N})\mathbf{A}\mathbf{q}_\nu \Big) \\ & + \frac{\lambda}{2N} \Big(\mathrm{tr}\{\mathbf{A}^\top \mathbf{K}(\mathbf{x}_{1:N}, \mathbf{x}_{1:N})\mathbf{A}\} + \sum_{\nu=1}^{N} \|\mathbf{q}_\nu\|_2^2 \Big) \end{aligned} \tag{5}$$

where the $N \times 1$ vector $\mathbf{k}(\mathbf{x}_{1:N}, \mathbf{x}_n)$ in (5) is the n-th column of $\mathbf{K}(\mathbf{x}_{1:N}, \mathbf{x}_{1:N})$, and since $\mathbf{A}$ has size $N \times r$, the minimization in (5) does not depend on $\bar{D}$.

Our goal is to develop and analyze offline as well as online solvers for (5). By pre-specifying an affordable complexity for the online solver, we aim at a low-complexity algorithm where subspace learning and feature extraction can be performed on-the-fly for streaming applications. Furthermore, we will introduce a novel approach to extracting features on which the kernel-based learning tasks of complexity $\mathcal{O}(N^3)$ can be well approximated by linear counterparts of complexity $\mathcal{O}(rN)$, hence realizing great savings in memory and computation while maintaining performance.

3 Offline kernel-based Feature Extraction via low-rank subspace tracking

Given a dataset $\{\mathbf{x}_\nu\}_{\nu=1}^N$ and leveraging the bi-convexity of the minimization in (5), we introduce in this section an offline solver, where two blocks of variables ($\mathbf{A}$ and $\{\mathbf{q}_\nu\}_{\nu=1}^N$) are updated alternately. The following two updates are carried out iteratively until convergence.

Update 1. With $\mathbf{A}[k]$ given from iteration k, the projection vectors $\{\mathbf{q}_\nu\}_{\nu=1}^N$ in iteration $k+1$ are updated as

$$\mathbf{q}_\nu[k+1] = \arg\min_{\mathbf{q}} \quad \ell(\mathbf{x}_\nu; \mathbf{A}[k], \mathbf{q}; \mathbf{x}_{1:N}) + \frac{\lambda}{2}\|\mathbf{q}\|_2^2 \quad (6a)$$

where the fitting cost $\ell(.)$ is given by [cf. (3)-(5)]

$$\begin{aligned}\ell(\mathbf{x}_\nu; \mathbf{A}[k], \mathbf{q}; \mathbf{x}_{1:N}) &:= \frac{1}{2}\|\boldsymbol{\phi}(\mathbf{x}_\nu) - \boldsymbol{\Phi}_N \mathbf{A}[k]\mathbf{q}\|_{\mathcal{H}}^2 \quad (6b)\\ &= \kappa(\mathbf{x}_\nu, \mathbf{x}_\nu) - 2\mathbf{k}^\top(\mathbf{x}_{1:N}, \mathbf{x}_\nu)\mathbf{A}[k]\mathbf{q} \\ &\quad + \mathbf{q}^\top \mathbf{A}^\top[k]\mathbf{K}(\mathbf{x}_{1:N}, \mathbf{x}_{1:N})\mathbf{A}[k]\mathbf{q}\,.\end{aligned}$$

The minimizer of (6a) yields the features as regularized projection coefficients of the lifted data vectors onto the virtual subspace $\bar{\mathbf{L}}_N[k] = \boldsymbol{\Phi}_N \mathbf{A}[k]$, and is given in closed form by

$$\begin{aligned}\mathbf{q}_\nu[k+1] = (\mathbf{A}^\top[k]\mathbf{K}(\mathbf{x}_{1:N}, \mathbf{x}_{1:N})\mathbf{A}[k] + \lambda \mathbf{I}_r)^{-1} \\ \times \mathbf{A}^\top[k]\mathbf{k}(\mathbf{x}_{1:N}, \mathbf{x}_\nu), \quad \nu = 1, ..., N\,. \quad (7)\end{aligned}$$

Update 2. With $\{\mathbf{q}_\nu[k+1]\}_{\nu=1}^N$ fixed and after dropping irrelevant terms, the subspace factor is obtained as [cf. (5)]

$$\begin{aligned}\mathbf{A}[k+1] = \arg\min_{\mathbf{A}} \frac{1}{N}\sum_{\nu=1}^N \ell(\mathbf{x}_\nu; \mathbf{A}, \mathbf{q}_\nu[k+1]; \mathbf{x}_{1:N}) \\ + \frac{\lambda}{2N}\mathrm{tr}\{\mathbf{A}^\top \mathbf{K}(\mathbf{x}_{1:N}, \mathbf{x}_{1:N})\mathbf{A}\}\,. \quad (8)\end{aligned}$$

Since $\mathbf{K}$ is strictly positive definite in practice, the task in (8) involves a strictly convex minimization. Equating the gradient to zero, yields the wanted subspace factor in closed form

$$\mathbf{A}[k+1] = \mathbf{Q}_N^\top[k+1]\Big(\mathbf{Q}_N[k+1]\mathbf{Q}_N^\top[k+1] + \lambda \mathbf{I}_r\Big)^{-1}\,. \quad (9)$$

Algorithm 1 provides the pseudocode for the update rules (7) and (9) of the offline solver, and the following proposition gives a guarantee on the convergence of the proposed solver to a local stationary point.

Proposition 1. *For positive definite kernels and $\lambda > 0$, the sequence $\{\mathbf{A}[k], \mathbf{Q}_N[k]\}$ generated by Algorithm 1 converges to a stationary point of the minimization in* (5).

Proof: Since the minimizations in (6a) and (8) are strictly convex with unique solutions, the result follows form [2, p. 272]. ■

Algorithm 1 BKFE: Batch Kernel-based Feature Extraction

Input $\{\mathbf{x}_\nu\}_{\nu=1}^N, \lambda$
Initialize $\mathbf{A}[1]$ at random
For $k = 1, \ldots$ do

$$\mathbf{S}[k+1] = \Big(\mathbf{A}^\top[k]\mathbf{K}(\mathbf{x}_{1:N}, \mathbf{x}_{1:N})\mathbf{A}[k] + \lambda \mathbf{I}_r\Big)^{-1}\mathbf{A}^\top[k]$$

$$\mathbf{Q}[k+1] = \mathbf{S}[k+1]\mathbf{K}(\mathbf{x}_{1:N}, \mathbf{x}_{1:N})$$

$$\mathbf{A}[k+1] = \mathbf{Q}_N^\top[k+1]\Big(\mathbf{Q}_N[k+1]\mathbf{Q}_N^\top[k+1] + \lambda \mathbf{I}_r\Big)^{-1}$$

Repeat Until Convergence
Return $\mathbf{A}[k], \{\mathbf{q}_\nu[k]\}_{\nu=1}^N$

Since matrix inversions in (7) and (9) cost $\mathcal{O}(r^3)$, and $\mathbf{Q}_N$ and $\mathbf{A}$ have size $r \times N$ and $N \times r$, respectively, the per iteration cost is $\mathcal{O}(N^2 r + N r^2 + r^3)$. Although the number of iterations needed in practice for Algorithm 1 to converge is effectively small, this per iteration complexity can be unaffordable for large datasets. In addition, datasets are not always available offline, or due to their massive volume, they can not be uploaded into memory at once. To cope with these issues, an online solver for (5) is developed next, where the updates are carried out by iteratively passing over the dataset one datum at a time.

4 Online kernel-based feature-extraction

This section deals with low-cost, on-the-fly updates of the 'virtual' subspace $\bar{\mathbf{L}}$, or equivalently its factor $\mathbf{A}$ as well as the features $\{\mathbf{q}_\nu\}$ that are desirable to keep up with streaming data. For such online updates, stochastic gradient descent (SGD) offers a practical approach, especially for *parametric* settings. However, upon processing n data vectors, $\mathbf{A}$ has size $n \times r$, which obviously grows with n. Hence, as the number of subspace factor parameters increases with the number of data, the task of interest is a *nonparametric* one. Unfortunately, performance of SGD on nonparametric learning such as the one at hand is a relatively uncharted territory. Nevertheless, SGD can still be performed on the initial formulation (2), where solving for the virtual $\bar{\mathbf{L}}$ constitutes a parametric task, not dependent on n.

Starting with an update for $\bar{\mathbf{L}}$, an update for $\mathbf{A}$ will be derived first, as an alternative to those in [8], [32], and [35]. Next, an SGD iteration for $\mathbf{A}$ will be developed in subsection 4.2, while in subsection 4.3 a connection between the two update rules will be drawn, suggesting how SGD can be broadened to learning nonparametric models as well.

4.1 SGD on "parametric" nonlinear subspace tracking

Suppose that $\mathbf{x}_n$ is acquired at time n, posing the overall joint subspace tracking and feature extraction problem as

[cf. (3)]

$$\min_{\bar{\mathbf{L}},\{\mathbf{q}_\nu\}_{\nu=1}^n} \frac{1}{2n}\sum_{\nu=1}^{n}\|\boldsymbol{\phi}(\mathbf{x}_\nu)-\bar{\mathbf{L}}\mathbf{q}_\nu\|_{\mathcal{H}}^2+\frac{\lambda}{2n}\Big(\|\bar{\mathbf{L}}\|_F^2+\|\mathbf{Q}_n\|_F^2\Big)\,. \tag{10}$$

Using an alternating minimization approach, we update features and the subspace per data vector as follows.

Update 1. Fixing the subspace estimate at its recent value $\bar{\mathbf{L}}[n-1]:=\boldsymbol{\Phi}_{n-1}\mathbf{A}[n-1]$ from time $n-1$, the projection vector of the new data vector $\mathbf{x}_n$ is found as [cf. (6a)]

$$\mathbf{q}[n]=\arg\min_{\mathbf{q}}\ell(\mathbf{x}_\nu;\mathbf{A}[n-1],\mathbf{q};\mathbf{x}_{1:n-1})+\frac{\lambda}{2}\|\mathbf{q}\|_2^2 \tag{11a}$$

which through the kernel trick readily yields

$$\begin{aligned}\mathbf{q}[n]=&(\mathbf{A}^\top[n-1]\mathbf{K}(\mathbf{x}_{1:n-1},\mathbf{x}_{1:n-1})\mathbf{A}[n-1]+\lambda\mathbf{I}_r)^{-1}\\&\times\mathbf{A}^\top[n-1]\mathbf{k}(\mathbf{x}_{1:n-1},\mathbf{x}_n)\,.\end{aligned} \tag{11b}$$

Although this update can be done for all the previous features $\{\mathbf{q}_\nu\}_{\nu=1}^{n-1}$ as well, it is skipped in practice to prevent exploding complexity. In the proposed algorithm, feature extraction is performed only for the most recent data vector $\mathbf{x}_n$.

Update 2. Having obtained $\mathbf{q}[n]$, the subspace update is given by solving

$$\min_{\bar{\mathbf{L}}}\frac{1}{n}\sum_{\nu=1}^{n}\bar{\mathcal{L}}(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q}[\nu]) \tag{12a}$$

where

$$\bar{\mathcal{L}}(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q}[\nu]):=\frac{1}{2}\|\boldsymbol{\phi}(\mathbf{x}_\nu)-\bar{\mathbf{L}}\mathbf{q}[\nu]\|_{\mathcal{H}}^2+\frac{\lambda}{2n}\|\bar{\mathbf{L}}\|_F^2\,. \tag{12b}$$

Solving (12a) as time evolves, becomes increasingly complex, and eventually unaffordable. If data $\{\mathbf{x}_\nu\}_{\nu=1}^n$ satisfy the law of large numbers, then (12a) approximates $\min_{\bar{\mathbf{L}}}\mathbb{E}[\bar{\mathcal{L}}(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q}_\nu)]$, where expectation is with respect to the unknown probability distribution of the data. To reduce complexity of the minimization, one typically resorts to stochastic approximation solvers, where by dropping the expectation (equivalently, the sample averaging operator), the 'virtual' subspace update is

$$\bar{\mathbf{L}}[n]=\bar{\mathbf{L}}[n-1]-\mu_{n,L}\bar{\mathbf{G}}_n \tag{13}$$

with $\mu_{n,L}$ denoting a preselected stepsize, and $\bar{\mathbf{G}}_n$ the gradient of the n-th summand in (12a) given by [cf. (12b)]

$$\begin{aligned}\bar{\mathbf{G}}_n&:=\nabla_{\bar{\mathbf{L}}}\bar{\mathcal{L}}(\mathbf{x}_n;\bar{\mathbf{L}}[n-1],\mathbf{q}[n])\\&=-\Big(\boldsymbol{\phi}(\mathbf{x}_n)-\bar{\mathbf{L}}[n-1]\mathbf{q}[n]\Big)\mathbf{q}^\top[n]+\frac{\lambda}{n}\bar{\mathbf{L}}[n-1]\\&=\boldsymbol{\Phi}_n\begin{bmatrix}\mathbf{A}[n-1]\mathbf{q}[n]\mathbf{q}^\top[n]\\-\mathbf{q}^\top[n]\end{bmatrix}+\frac{\lambda}{n}\boldsymbol{\Phi}_n\begin{bmatrix}\mathbf{A}[n-1]\\\mathbf{0}_{1\times r}\end{bmatrix}\,.\end{aligned} \tag{14}$$

Using (4) to rewrite $\bar{\mathbf{L}}[n]=\boldsymbol{\Phi}_n\mathbf{A}[n]$, and substituting into (13), yields

$$\begin{aligned}\boldsymbol{\Phi}_n\mathbf{A}[n]=&\boldsymbol{\Phi}_n\begin{bmatrix}\mathbf{A}[n-1]\\\mathbf{0}_{1\times r}\end{bmatrix}\\&-\mu_{n,L}\boldsymbol{\Phi}_n\begin{bmatrix}\mathbf{A}[n-1]\Big(\mathbf{q}[n]\mathbf{q}^\top[n]+\frac{\lambda}{n}\mathbf{I}_r\Big)\\-\mathbf{q}^\top[n]\end{bmatrix}\end{aligned} \tag{15}$$

which suggests the following update rule for the subspace factor

$$\mathbf{A}[n]=\begin{bmatrix}\mathbf{A}[n-1]-\mu_{n,L}\mathbf{A}[n-1]\Big(\mathbf{q}[n]\mathbf{q}^\top[n]+\frac{\lambda}{n}\mathbf{I}_r\Big)\\\mu_{n,L}\mathbf{q}^\top[n]\end{bmatrix}\,. \tag{16}$$

Even though (16) is not the only iteration satisfying (15), it offers an efficient update of the factor $\mathbf{A}$. The update steps for the proposed parametric tracker are summarized as Algorithm 2.

Note that the multiplication and inversion in (9) are avoided. However, per data vector processed, the kernel matrix is expanded by one row and one column, while the subspace factor $\mathbf{A}$ grows accordingly by one row.

Algorithm 2 Online kernel-based feature extraction (OK-FE)

Input $\{\mathbf{x}_\nu\}_{\nu=1}^n,\lambda$
Initialize $\mathbf{A}[1]=\mathbf{1}_{1\times r}$, $\mathbf{K}(\mathbf{x}_1,\mathbf{x}_1)=\kappa(\mathbf{x}_1,\mathbf{x}_1)$
For $n=2,\ldots$ do

$$\begin{aligned}\mathbf{q}[n]=&(\mathbf{A}^\top[n-1]\mathbf{K}(\mathbf{x}_{1:n-1},\mathbf{x}_{1:n-1})\mathbf{A}[n-1]+\lambda\mathbf{I}_r)^{-1}\\&\times\mathbf{A}^\top[n-1]\mathbf{k}(\mathbf{x}_{1:n-1},\mathbf{x}_n)\end{aligned}$$

$$\mathbf{K}(\mathbf{x}_{1:n},\mathbf{x}_{1:n})=\begin{bmatrix}\mathbf{K}(\mathbf{x}_{1:n-1},\mathbf{x}_{1:n-1})&\mathbf{k}(\mathbf{x}_{1:n-1},\mathbf{x}_n)\\\mathbf{k}^\top(\mathbf{x}_{1:n-1},\mathbf{x}_n)&\kappa(\mathbf{x}_n,\mathbf{x}_n)\end{bmatrix}$$

$$\mathbf{A}[n]=\begin{bmatrix}\mathbf{A}[n-1]-\mu_{n,L}\mathbf{A}[n-1]\Big(\mathbf{q}[n]\mathbf{q}^\top[n]+\frac{\lambda}{n}\mathbf{I}_r\Big)\\\mu_{n,L}\mathbf{q}^\top[n]\end{bmatrix}$$

Return $\mathbf{A}[n],\{\mathbf{q}[\nu]\}_{\nu=2}^n$

4.2 SGD for nonparametric subspace tracking

In this subsection, the feature extraction rule in (11b) is retained, while the update rule (16) is replaced by directly acquiring the SGD direction along the gradient of the instantaneous objective term with respect to $\mathbf{A}$. Since, in contrast to the fixed-size matrix $\bar{\mathbf{L}}$, the number of parameters in $\mathbf{A}$ grows with n, we refer to the solver developed in this subsection as *nonparametric* subspace tracker. Furthermore, the connection between the two solvers is drawn in subsection 4.3, and convergence of the proposed algorithm is analyzed in subsection 4.4.

At time instance n, subproblem (12a) can be expanded using the kernel trick as

$$\min_{\mathbf{A}\in\mathbb{R}^{n\times r}}\frac{1}{n}\sum_{\nu=1}^{n}\mathcal{L}(\mathbf{x}_\nu;\mathbf{A},\mathbf{q}[\nu];\mathbf{x}_{1:n})\} \tag{17}$$

where

$$\begin{aligned}\mathcal{L}(\mathbf{x}_\nu;\mathbf{A},\mathbf{q}[\nu];\mathbf{x}_{1:n}):=&\,\ell(\mathbf{x}_\nu;\mathbf{A},\mathbf{q}[\nu];\mathbf{x}_{1:n})\\&+\frac{\lambda}{2n}\mathrm{tr}\{\mathbf{A}^\top\mathbf{K}(\mathbf{x}_{1:n},\mathbf{x}_{1:n})\mathbf{A}\}\end{aligned} \tag{18}$$

with $\ell(.)$ given by (6b). Stochastic approximation solvers of (17) suggest the update

$$\mathbf{A}[n]=\begin{bmatrix}\mathbf{A}[n-1]\\\mathbf{0}_{r\times 1}^\top\end{bmatrix}-\mu_{n,A}\mathbf{G}_n \tag{19a}$$

where $\mu_{n,A}$ denotes the user-selected step size, and $\mathbf{G}_n$ denotes the gradient of the n-th summand in (17) with respect to $\mathbf{A}$ that is given by

$$\begin{aligned}\mathbf{G}_n :=& \nabla_{\mathbf{A}}\mathcal{L}(\mathbf{x}_n;[\mathbf{A}^\top[n-1],\mathbf{0}_{r\times 1}]^\top,\mathbf{q}[n];\mathbf{x}_{1:n}) \qquad (19b)\\ =& \mathbf{K}(\mathbf{x}_{1:n},\mathbf{x}_{1:n})\begin{bmatrix}\mathbf{A}[n-1]\\ \mathbf{0}_{r\times 1}^\top\end{bmatrix}\mathbf{q}[n]\mathbf{q}^\top[n]\\ &-\mathbf{k}(\mathbf{x}_{1:n},\mathbf{x}_n)\mathbf{q}^\top[n]+\frac{\lambda}{n}\mathbf{K}(\mathbf{x}_{1:n},\mathbf{x}_{1:n})\begin{bmatrix}\mathbf{A}[n-1]\\ \mathbf{0}_{r\times 1}^\top\end{bmatrix}.\end{aligned}$$

Substituting (19b) into (19a) yields the desired update of $\mathbf{A}$ which together with (11b) constitute our nonparametroc solver, tabulated under Algorithm 3.

Algorithm 3 Online kernel-based feature extraction (OK-FE)

Input $\{\mathbf{x}_\nu\}_{\nu=1}^n, \lambda$
Initialize $\mathbf{A}[1]=\mathbf{1}_{1\times r}$, $\mathbf{K}(\mathbf{x}_1,\mathbf{x}_1)=\kappa(\mathbf{x}_1,\mathbf{x}_1)$
For $n=2,\ldots$ do

$$\mathbf{q}[n]=(\mathbf{A}^\top[n-1]\mathbf{K}(\mathbf{x}_{1:n-1},\mathbf{x}_{1:n-1})\mathbf{A}[n-1]+\lambda\mathbf{I}_r)^{-1}\times\mathbf{A}^\top[n-1]\mathbf{k}(\mathbf{x}_{1:n-1},\mathbf{x}_n)$$

$$\mathbf{K}(\mathbf{x}_{1:n},\mathbf{x}_{1:n})=\begin{bmatrix}\mathbf{K}(\mathbf{x}_{1:n-1},\mathbf{x}_{1:n-1}) & \mathbf{k}(\mathbf{x}_{1:n-1},\mathbf{x}_n)\\ \mathbf{k}^\top(\mathbf{x}_{1:n-1},\mathbf{x}_n) & \kappa(\mathbf{x}_n,\mathbf{x}_n)\end{bmatrix}$$

$$\mathbf{G}_n=\mathbf{K}(\mathbf{x}_{1:n},\mathbf{x}_{1:n})\begin{bmatrix}\mathbf{A}[n-1]\\ \mathbf{0}_{r\times 1}^\top\end{bmatrix}\mathbf{q}[n]\mathbf{q}^\top[n]-\mathbf{k}(\mathbf{x}_{1:n},\mathbf{x}_n)\mathbf{q}^\top[n]+\frac{\lambda}{n}\mathbf{K}(\mathbf{x}_{1:n},\mathbf{x}_{1:n})\begin{bmatrix}\mathbf{A}[n-1]\\ \mathbf{0}_{r\times 1}^\top\end{bmatrix}$$

$$\mathbf{A}[n]=\begin{bmatrix}\mathbf{A}[n-1]\\ \mathbf{0}_{r\times 1}^\top\end{bmatrix}-\mu_{n,A}\mathbf{G}_n$$

Return $\mathbf{A}[n],\{\mathbf{q}[\nu]\}_{\nu=2}^n$

4.3 Linking parametric with nonparametric SGD updates

Considering that $\bar{\mathbf{L}}[n]=\mathbf{\Phi}_n\mathbf{A}[n]$ holds for all n, it is apparent from (19b) and (14) that $\mathbf{G}_n=\mathbf{\Phi}_n^\top\bar{\mathbf{G}}_n$. The latter implies that the update rule in (19a) amounts to performing SGD on $\bar{\mathbf{L}}$ with a matrix stepsize $\mathbf{D}_n=\mathbf{\Phi}_n\mathbf{\Phi}_n^\top$; that is,

$$\bar{\mathbf{L}}[n]=\bar{\mathbf{L}}[n-1]-\mu_{n,A}\mathbf{D}_n\bar{\mathbf{G}}_n\,. \qquad (20)$$

Consequently, it is important to note that such a choice constitutes a valid descent direction, which is guaranteed since

$$\bar{\mathbf{G}}_n^\top\mathbf{D}_n\bar{\mathbf{G}}_n=\mathbf{H}_n^\top\mathbf{K}^\top(\mathbf{x}_{1:n},\mathbf{x}_{1:n})\mathbf{K}(\mathbf{x}_{1:n},\mathbf{x}_{1:n})\mathbf{H}_n\succcurlyeq\mathbf{0} \qquad (21)$$

where

$$\mathbf{H}_n:=\begin{bmatrix}\mathbf{A}[n-1](\mathbf{q}_n\mathbf{q}_n^\top+\frac{\lambda}{n}\mathbf{I}_r)\\ -\mathbf{q}_n^\top\end{bmatrix}.$$

Moreover, for positive definite kernel matrices, formed by e.g., Gaussian kernels, we have $\bar{\mathbf{G}}_n^\top\mathbf{D}_n\bar{\mathbf{G}}_n\succ\mathbf{0}$, which strongly guarantees that $-\mathbf{D}_n\bar{\mathbf{G}}_n$ is a descent direction [2, p. 35]. Leveraging this link, Algorithm 3 will be shown next to enjoy the same convergence guarantee as that of Algorithm 2.

Remark 1. Although SGD solver in Algorithm 3 can be viewed as a special case of Algorithm 2, developing the parametric SGD solver in Algorithm 2 will allow us to analyze convergence of the two algorithms next.

4.4 Convergence analysis

The cost in (10) can be written as

$$F_n(\bar{\mathbf{L}}):=\frac{1}{n}\sum_{\nu=1}^n\min_{\mathbf{q}}f_\nu(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q}) \qquad (22)$$

with $f_\nu(\mathbf{x}_\nu,\bar{\mathbf{L}},\mathbf{q}):=\bar{\mathcal{L}}(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q})+(\lambda/2)\|\mathbf{q}\|_2^2$, and $\bar{\mathcal{L}}$ as in (12b). Thus, the minimization in (10) is equivalent to $\min_{\bar{\mathbf{L}}}F_n(\bar{\mathbf{L}})$. To ensure convergence of the proposed algorithms, the following assumptions are adopted.

(A1) *Data vectors are drawn independently from an identical distribution; and*

(A2) *The sequence* $\{\|\bar{\mathbf{L}}[\nu]\|_F\}_{\nu=1}^\infty$ *is bounded.*

The independence of $\{\mathbf{x}_\nu\}$ across time is standard when studying the performance of online algorithms [24], while boundedness of the iterates $\{\|\bar{\mathbf{L}}[\nu]\|_F\}_{\nu=1}^\infty$ is a technical condition that simplifies the analysis, and in the present setting is guaranteed due to the Frobenious-norm regularization (by proper choice of the stepsize).

Proposition 2. *Under (A1)-(A2), if* $\mu_{n,L}=1/\bar{\gamma}_n$ *with* $\bar{\gamma}_n:=\sum_{\nu=1}^n\gamma_\nu$ *and* $\gamma_\nu\geq\|\nabla^2\bar{\mathcal{L}}(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q}[\nu])\|_{\mathcal{H}}\ \ \forall n$, *then the subspace iterates in* (13) *satisfy* $\lim_{n\to\infty}\nabla F_n(\bar{\mathbf{L}}[n])=\mathbf{0}$ *almost surely; that is, they fall into the stationary point of* (10).

Proof: The proof of the proposition is inspired by [24] and [23], and can be sketched along the following steps.

Step 1. First, we judiciously introduce a surrogate for $F_n(\bar{\mathbf{L}})$ whose minimizer coincides with the SGD updates in (13).

To this end, we have that $\min_{\mathbf{q}}f_\nu(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q})\leq f_\nu(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q}[\nu])$; hence, $\hat{F}_n(\bar{\mathbf{L}}):=(1/n)\sum_{\nu=1}^n f_\nu(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q}[\nu])$ upper bounds the cost function, namely $F_n(\bar{\mathbf{L}})\leq\hat{F}_n(\bar{\mathbf{L}}),\ \forall\bar{\mathbf{L}}$. Further approximating f_n through a second-order Taylor's expansion at the previous subspace update $\bar{\mathbf{L}}[n-1]$, we arrive at

$$\begin{aligned}\tilde{f}_n(\mathbf{x}_n;\bar{\mathbf{L}},\mathbf{q}[n])=& f_n(\mathbf{x}_n;\bar{\mathbf{L}}[n-1],\mathbf{q}[n]) \qquad (23)\\ &+\mathrm{tr}\{\nabla_{\bar{\mathbf{L}}}f_n(\mathbf{x}_n;\bar{\mathbf{L}}[n-1],\mathbf{q}[n])(\bar{\mathbf{L}}-\bar{\mathbf{L}}[n-1])^\top\}\\ &+\frac{\gamma_n}{2}\|\bar{\mathbf{L}}-\bar{\mathbf{L}}[n-1]\|_F^2\,.\end{aligned}$$

By choosing $\gamma_n\geq\|\nabla^2_{\bar{\mathbf{L}}}f_n(\mathbf{x}_n;\bar{\mathbf{L}}[n-1],\mathbf{q}_n)\|_{\mathcal{H}}=\|(\mathbf{q}[n]\mathbf{q}^\top[n])\otimes\mathbf{I}_{\bar{D}}+(\lambda/n)\mathbf{I}_{r\bar{D}}\|_{\mathcal{H}}$ and using the norm properties in the Hilbert space, it is important to recognize that:

(i) $\tilde{f}_n$ is locally tight; i.e., $\tilde{f}_n(\mathbf{x}_n;\bar{\mathbf{L}}[n-1],\mathbf{q}[n])=f_n(\mathbf{x}_n;\bar{\mathbf{L}}[n-1],\mathbf{q}[n])$;

(ii) gradient of $\tilde{f}_n$ is locally tight; i.e., $\nabla_{\bar{\mathbf{L}}}\tilde{f}_n(\mathbf{x}_n;\bar{\mathbf{L}}[n-1],\mathbf{q}[n])=\nabla_{\bar{\mathbf{L}}}f_n(\mathbf{x}_n;\bar{\mathbf{L}}[n-1],\mathbf{q}[n])$; and

(iii) $\tilde{f}_n$ *globally* majorizes the original instantaneous cost f_n; that is, $f_n(\mathbf{x}_n;\bar{\mathbf{L}},\mathbf{q}[n])\leq\tilde{f}_n(\mathbf{x}_n;\bar{\mathbf{L}},\mathbf{q}[n]),\forall\bar{\mathbf{L}}$.

Selecting now the target surrogate cost as

$$\tilde{F}_n(\bar{\mathbf{L}})=\frac{1}{n}\sum_{\nu=1}^n\tilde{f}_\nu(\mathbf{x}_\nu;\bar{\mathbf{L}},\mathbf{q}[\nu])$$

we have $F_n(\bar{\mathbf{L}})\leq\hat{F}_n(\bar{\mathbf{L}})\leq\tilde{F}_n(\bar{\mathbf{L}}),\forall\bar{\mathbf{L}}$. Minimizing the cost $\tilde{F}_n(\bar{\mathbf{L}})$ amounts to nullifying the gradient, i.e., $\nabla_{\bar{\mathbf{L}}}\tilde{F}_n(\bar{\mathbf{L}}[n])=\mathbf{0}$, which yields [24]

$$\bar{\mathbf{L}}[n]=\bar{\mathbf{L}}[n-1]-\frac{1}{\bar{\gamma}_n}\bar{\mathbf{G}}_n \qquad (24)$$

where $\bar{\gamma}_n := \sum_{\nu=1}^{n} \gamma_\nu$. By setting $\mu_n = 1/\bar{\gamma}_n$, the SGD-based update of $\bar{\mathbf{L}}[n]$ now coincides with the minimizer of $\tilde{F}_n(\bar{\mathbf{L}})$; that is, $\bar{\mathbf{L}}[n] = \arg\min_{\bar{\mathbf{L}}} \tilde{F}_n(\bar{\mathbf{L}})$.

Step 2. The second step establishes that the surrogate costs $\{\tilde{F}_n(\bar{\mathbf{L}})\}$ form a quasi-martingale sequence, and using tightness of the surrogate cost we deduce that $\lim_{n\to\infty}(F_n(\bar{\mathbf{L}}[n]) - \tilde{F}_n(\bar{\mathbf{L}}[n])) = 0$. Thus, the surrogate cost asymptotically converges to the original cost $F_n(\bar{\mathbf{L}})$.

Step 3. Leveraging the regularity of $\tilde{\mathcal{L}}(\mathbf{x}_\nu; \bar{\mathbf{L}}, \mathbf{q}_\nu)$, convergence of the cost sequence implies convergence of $\{\|\nabla_{\bar{\mathbf{L}}} F_n(\bar{\mathbf{L}}[n]) - \nabla_{\bar{\mathbf{L}}} \tilde{F}_n(\bar{\mathbf{L}}[n])\|_{\mathcal{H}}\}$ to zero, which along with $\nabla_{\bar{\mathbf{L}}} \tilde{F}_n(\bar{\mathbf{L}}[n]) = \mathbf{0}$, yields $\{\|\nabla_{\bar{\mathbf{L}}} F_n(\bar{\mathbf{L}}[n])\|_{\mathcal{H}}\} \to \mathbf{0}$. ■

So far, we have asserted convergence of the SGD-based algorithm for the "virtual" $\bar{\mathbf{L}}$ provided by Algorithm 2. A related convergence result for Algorithm 3 is guaranteed by the following argument.

Proposition 3. *Under (A1)-(A2) and for positive definite kernels, if $\mu_{n,A} = 1/\bar{\xi}_n$ with $\bar{\xi}_n := \sum_{\nu=1}^{n} \xi_n$ and $\xi_n \geq n\gamma_n$, then the subspace iterates in* (19a) *satisfy $\lim_{n\to\infty} \nabla C_n(\bar{\mathbf{L}}[n]) = \mathbf{0}$ almost surely; that is, the subspace iterates will converge to the stationary point of* (10).

Proof: The proof follows the steps in Proposition 2, with an extra step in the construction of the appropriate surrogate cost in Step 1. In particular, using that $\forall n$ the optimal subspace is of the form $\bar{\mathbf{L}}_n = \mathbf{\Phi}_n \mathbf{A}$, the objective $\tilde{f}_\nu$ can be further majorized over the subset of virtual subspaces $\bar{\mathbf{L}} = \mathbf{\Phi}_n \mathbf{A}$, by

$$\begin{aligned}\check{f}_n(\mathbf{x}_n; \mathbf{\Phi}_n, \mathbf{A}, \mathbf{q}[n]) &:= f_n(\mathbf{x}_n; \bar{\mathbf{L}}[n-1], \mathbf{q}[n]) \\ &+ \mathrm{tr}\{\nabla_{\bar{\mathbf{L}}} f_n(\mathbf{x}_n; \bar{\mathbf{L}}[n-1], \mathbf{q}[n])(\mathbf{\Phi}_n \mathbf{A} - \bar{\mathbf{L}}[n-1])^\top\} \\ &+ \frac{\xi_n}{2}\|\mathbf{A} - \begin{bmatrix}\mathbf{A}[n-1] \\ \mathbf{0}_{1\times r}\end{bmatrix}\|_F^2\end{aligned}$$

for which we have

$$\begin{aligned}\tilde{f}_n(\mathbf{x}_n; \bar{\mathbf{L}}, \mathbf{q}[n]) - \check{f}_\nu(\mathbf{x}_\nu; \mathbf{\Phi}_n, \mathbf{A}, \mathbf{q}_\nu) \\ = \frac{\gamma_n}{2}\|\bar{\mathbf{L}} - \bar{\mathbf{L}}[n-1]\|_F^2 - \frac{\xi_n}{2}\|\mathbf{A} - \begin{bmatrix}\mathbf{A}[n-1] \\ \mathbf{0}_{1\times r}\end{bmatrix}\|_F^2.\end{aligned}$$

The Cauchy-Schwarz inequality implies that

$$\begin{aligned}\|\bar{\mathbf{L}} - \bar{\mathbf{L}}[n-1]\|_F^2 &= \|\mathbf{\Phi}_n \mathbf{A} - \mathbf{\Phi}_n \begin{bmatrix}\mathbf{A}[n-1] \\ \mathbf{0}_{1\times r}\end{bmatrix}\|_F^2 \\ &\leq \|\mathbf{\Phi}_n\|_F^2 \|\mathbf{A} - \begin{bmatrix}\mathbf{A}[n-1] \\ \mathbf{0}_{1\times r}\end{bmatrix}\|_F^2\end{aligned}$$

and by choosing $\xi_n \geq \|\mathbf{\Phi}_n\|_F^2 \gamma_n = n\gamma_n$, we will have $\tilde{f}_n(\mathbf{x}_n; \bar{\mathbf{L}}, \mathbf{q}[n]) \leq \check{f}_\nu(\mathbf{x}_\nu; \mathbf{\Phi}_n, \mathbf{A}, \mathbf{q}_\nu)$. Selecting now $\check{f}_\nu(.)$ as the new surrogate whose minimizer coincides with the update rule in (19a), the rest of the proof follows that of Proposition 2. ■

5 Reduced-complexity OK-FE on a budget

Per data vector processed, the iterative solvers of the previous section have one column of $\mathbf{\Phi}_n$ and one row of $\mathbf{A}$ added, which implies growing memory and complexity requirements as n grows. The present section combines two means of coping with this formidable challenge: one based on *censoring* uninformative data, and the second based on "budget maintenance." By modifying Algorithms 2 and 3 accordingly, memory and complexity requirements are rendered affordable.

5.1 Censoring uninformative data

In the LS cost that Algorithms 2 and 3 rely on, small values of the fitting error can be tolerated in practice without noticeable performance degradation. This suggests modifying the LS cost so that small fitting errors (say up to $\pm\epsilon$) induce no penalty, e.g., by invoking the $\epsilon-$insensitive cost that is popular in support vector regression (SVR) settings [15].

Consider henceforth positive-definite kernels for which low-rank factors offer an approximation to the full-rank kernel matrix, and lead to a generally nonzero LS-fit $\|\mathbf{\Phi}_n - \bar{\mathbf{L}}\mathbf{Q}_n\|_{\mathcal{H}}$. These considerations suggest modifying the LS cost $\ell(\mathbf{x}_n; \mathbf{A}[n-1], \mathbf{q}; \mathbf{x}_{1:n-1})$ as

$$\begin{aligned}&\check{\ell}(\mathbf{x}_n; \mathbf{A}[n-1], \mathbf{q}; \mathbf{x}_{1:n-1}) \\ &:= \begin{cases} 0 & \text{if } \ell(\mathbf{x}_n; \mathbf{A}[n-1], \mathbf{q}; \mathbf{x}_{1:n-1}) < \epsilon \\ \ell(\mathbf{x}_n; \mathbf{A}[n-1], \mathbf{q}; \mathbf{x}_{1:n-1}) - \epsilon & \text{otherwise.}\end{cases}\end{aligned} \tag{25}$$

By proper choice of ϵ, the cost $\check{\ell}(.)$ implies that if $\ell(\mathbf{x}_n; \mathbf{A}[n-1], \mathbf{q}_n; \mathbf{x}_{1:n-1}) < \epsilon$, the virtual datum $\boldsymbol{\phi}(\mathbf{x}_n)$ is captured well enough by the virtual current subspace $\bar{\mathbf{L}}[n-1] = \mathbf{\Phi}_{n-1}\mathbf{A}[n-1]$, and the solver will not attempt to decrease its LS error, which suggests skipping the augmentation of the basis $\mathbf{\Phi}_{n-1}$, with the new lifted datum $\boldsymbol{\phi}(\mathbf{x}_n)$ [3].

In short, if the upper branch of (25) is in effect, $\boldsymbol{\phi}(\mathbf{x}_n)$ is considered uninformative, and it is censored for the subspace update step; whereas having the lower branch deems $\boldsymbol{\phi}(\mathbf{x}_n)$ informative, and augments the basis set of the virtual subspace. The latter case gives rise to what we term *online support vectors* (OSV), which must be stored, unlike 'censored' data that are discarded from subsequent subspace updates.

In order to keep track of the OSVs, let $\mathcal{S}_{n-1}$ denote the set of indices corresponding to the SVs revealed up to time n. Accordingly, rewrite $\bar{\mathbf{L}}[n-1] = \mathbf{\Phi}_{\mathcal{S}_{n-1}}\mathbf{A}[n-1]$, and the modified LS cost as $\check{\ell}(\mathbf{x}_n; \mathbf{A}[n-1], \mathbf{q}; \mathbf{x}_{\mathcal{S}_{n-1}})$, depending on which of the following two cases emerges.

C1. *If $\check{\ell}(\mathbf{x}_n; \mathbf{A}[n-1], \mathbf{q}; \mathbf{x}_{\mathcal{S}_{n-1}}) \leq \epsilon$, the OSV set will not grow, and we will have $\mathcal{S}_n = \mathcal{S}_{n-1}$; or,*

C2. *If $\check{\ell}(\mathbf{x}_n; \mathbf{A}[n-1], \mathbf{q}; \mathbf{x}_{\mathcal{S}_{n-1}}) > \epsilon$, the OSV set will grow, and we will have $\mathcal{S}_n = \mathcal{S}_{n-1} \cup \{n\}$.*

The subspace matrix per iteration will thus take the form $\bar{\mathbf{L}}[n] = \mathbf{\Phi}_{\mathcal{S}_n}\mathbf{A}[n]$, where $\mathbf{\Phi}_{\mathcal{S}_n} := [\boldsymbol{\phi}_{n_1}, ..., \boldsymbol{\phi}_{n_{|\mathcal{S}_n|}}]$, with $\mathcal{S}_n := \{n_1, n_2, ..., n_{|\mathcal{S}_n|}\}$, and $\mathbf{A} \in \mathbb{R}^{|\mathcal{S}_n|\times r}$. Upon replacing $\mathbf{x}_{1:n}$ in Algorithm 3 with $\mathbf{x}_{\mathcal{S}_n}$, Algorithm 4 gives the pseudocode for our reduced-complexity OKFE, which also includes a budget maintenance module that will be presented in the ensuing Section 5.2.

Modifying the LS-fit in (25) and discarding the censored data, certainly reduce the rate at which the memory and complexity requirements increase. Albeit at a slower rate, $|\mathcal{S}_n|$ may still grow unbounded as time proceeds. Thus, one is motivated to restrict the number of OSVs to a prescribed affordable budget, $|\mathcal{S}_n| \leq B$, and introduce a solver which maintains such a budget throughout the iterations. To this end, we introduce next a greedy 'budget maintenance' scheme.

5.2 Budget maintenance

When inclusion of a new data vector into the OSV set pushes its cardinality $|\mathcal{S}_n|$ beyond the prescribed budget B, the budget maintenance module will discard the SV whose exclusion distorts minimally $\bar{\mathbf{L}}[n]$. Specifically, with $\mathbf{\Phi}_{n\backslash i}$ and $\mathbf{A}_{\backslash i}[n]$ denoting $\mathbf{\Phi}_n$ and $\mathbf{A}[n]$ devoid of their i-th column and row, respectively, our rule for selecting the index to be excluded is

$$\begin{aligned} i_* &= \arg\min_{i\in\mathcal{S}_n} \|\mathbf{\Phi}_n\mathbf{A}[n] - \mathbf{\Phi}_{n\backslash i}\mathbf{A}_{\backslash i}[n]\|_F^2 \\ &= \arg\min_{i\in\mathcal{S}_n} \mathrm{tr}\{\mathbf{A}^\top[n]\mathbf{K}(\mathbf{x}_{\mathcal{S}_n},\mathbf{x}_{\mathcal{S}_n})\mathbf{A}[n] \\ &\qquad - 2\mathbf{A}_{\backslash i}^\top[n]\mathbf{K}(\mathbf{x}_{\mathcal{S}_n\backslash i},\mathbf{x}_{\mathcal{S}_n})\mathbf{A}[n] \\ &\qquad + \mathbf{A}_{\backslash i}^\top[n]\mathbf{K}(\mathbf{x}_{\mathcal{S}_n\backslash i},\mathbf{x}_{\mathcal{S}_n\backslash i})\mathbf{A}_{\backslash i}[n]\}\,. \end{aligned} \tag{26}$$

Enumeration over $\mathcal{S}_n$ and evaluation of the cost incurs complexity $\mathcal{O}(B^3)$ for solving (26). Hence, in order to mitigate the computational complexity, a greedy scheme is put forth. Since exclusion of an SV will result in removing the corresponding row from the subspace factor, discarding the SV corresponding to the row with the smallest ℓ_2-norm suggests a reasonable heuristic greedy policy. To this end, one needs to find the index

$$\hat{i}_* = \arg\min_{i=1,2,\ldots,B+1} \|\mathbf{a}_i[n]\|_2 \tag{27}$$

where $\mathbf{a}_i^\top[n]$ denotes the $i-$th row of $\mathbf{A}[n]$. Subsequently, $\hat{i}_*$ as well as the corresponding SV are discarded from $\mathcal{S}_n$ and the SV set respectively, and an OSV set of cardinality $|\mathcal{S}_n| = B$ is maintained. Algorithms 4 and 5 tabulate the updates and the proposed greedy budget maintenance scheme, respectively.

Remark 2. In principle, methods related to those in [35], including replacement of two SVs by a linear combination of the two, or projecting an SV on the SV set and discarding the projected SV, are also viable alternatives. In practice however, their improved performance relative to (27) is negligible and along with their increased complexity, renders such alternatives less attractive for large-scale datasets.

Algorithm 4 OK-FEB Algorithm

Input $\{\mathbf{x}_\nu\}_{\nu=1}^n, \lambda$
Initialize $\mathbf{A}[1]$ at random and $\mathcal{S}_1 = \{1\}$
For $n = 2, \ldots$ do

$$\begin{aligned} \mathbf{q}[n] =& (\mathbf{A}^\top[n-1]\mathbf{K}(\mathbf{x}_{\mathcal{S}_{n-1}},\mathbf{x}_{\mathcal{S}_{n-1}})\mathbf{A}[n-1] + \lambda\mathbf{I}_r)^{-1} \\ &\times \mathbf{A}^\top[n-1]\mathbf{k}(\mathbf{x}_{\mathcal{S}_{n-1}},\mathbf{x}_n) \\ \ell_n =& k(\mathbf{x}_n,\mathbf{x}_n) - 2\mathbf{k}^\top(\mathbf{x}_{\mathcal{S}_{n-1}},\mathbf{x}_n)\mathbf{A}[n-1]\mathbf{q}[n] \\ &+ \mathbf{q}_n^\top\mathbf{A}^\top[n-1]\mathbf{K}(\mathbf{x}_{\mathcal{S}_{n-1}},\mathbf{x}_{\mathcal{S}_{n-1}})\mathbf{A}[n-1]\mathbf{q}[n] \end{aligned}$$

if $\ell_n < \epsilon$ **then**
  $\mathcal{S}_n = \mathcal{S}_{n-1}$
else
  $\mathcal{S}_n = \mathcal{S}_{n-1} \cup \{n\}$

$$\begin{aligned} \check{\mathbf{G}}_n =& \mathbf{K}(\mathbf{x}_{\mathcal{S}_n},\mathbf{x}_{\mathcal{S}_n}) \begin{bmatrix}\mathbf{A}[n-1] \\ \mathbf{0}_{r\times 1}^\top\end{bmatrix} \mathbf{q}[n]\mathbf{q}^\top[n] \\ &- \mathbf{k}(\mathbf{x}_{\mathcal{S}_n},\mathbf{x}_n)\mathbf{q}^\top[n] + \frac{\lambda}{n}\mathbf{K}(\mathbf{x}_{\mathcal{S}_n},\mathbf{x}_{\mathcal{S}_n}) \begin{bmatrix}\mathbf{A}[n-1] \\ \mathbf{0}_{r\times 1}^\top\end{bmatrix} \end{aligned}$$

  $\mathbf{A}[n] = \begin{bmatrix}\mathbf{A}[n-1] \\ \mathbf{0}_{r\times 1}^\top\end{bmatrix} - \mu_{n,A}\check{\mathbf{G}}_n$
  if $|\mathcal{S}_n| > B$ **then**
    Run budget maintenance module
  end if
end if
EndFor
Return $\mathbf{A}[n], \mathcal{S}_n, \{\mathbf{q}[\nu]\}_{\nu=1}^n$

Algorithm 5 Budget maintenace module

Input $\{\mathcal{S}, \mathbf{A}\}$
$\hat{i}_* = \arg\min_{i\in\mathcal{S}} \|\mathbf{a}_i^\top\|_2$
$\mathcal{S} \leftarrow \mathcal{S} \setminus \{\hat{i}_*\}$
Remove the $\hat{i}_*$-th row of $\mathbf{A}$
Return $\{\mathcal{S}, \mathbf{A}\}$

5.3 Complexity analysis

Computational complexity of the proposed OK-FEB algorithm is evaluated in the present section. The computations required by the $n-$th iteration of Algorithm 4 for feature extraction and parameter update depend on B, r, and D, as well as the censoring process outlined in Section 5.1. Specifically, computing $\check{\mathbf{G}}_n$ and performing the first-order stochastic update that yields $\mathbf{A}[n]$ requires $\mathbf{O}(B^2r)$ multiplications, a cost that is saved for iterations that are skipped when $\ell_n < \epsilon$. Regarding the computation of $\mathbf{q}[n]$, $Br(B+r)$ multiplications are needed to form $\mathbf{A}^\top[n-1]\mathbf{K}(\mathbf{x}_{\mathcal{S}_{n-1}},\mathbf{x}_{\mathcal{S}_{n-1}})\mathbf{A}[n-1]$, and $\mathcal{O}(r^3)$ multiplications for the inversion of $\mathbf{A}^\top[n-1]\mathbf{K}(\mathbf{x}_{\mathcal{S}_{n-1}},\mathbf{x}_{\mathcal{S}_{n-1}})\mathbf{A}[n-1] + \lambda\mathbf{I}_r$. Fortunately, the aforementioned computations can also be avoided for iteration n, if the previous iteration $n-1$ performs no update on $\mathbf{A}[n-1]$; in this case, $(\mathbf{A}^\top[n-1]\mathbf{K}(\mathbf{x}_{\mathcal{S}_{n-1}},\mathbf{x}_{\mathcal{S}_{n-1}})\mathbf{A}[n-1] + \lambda\mathbf{I}_r)^{-1}$ remains unchanged and can simply be accessed from memory. Nevertheless, a "baseline" of computations is required for feature extraction related operations that take place regardless of censoring. Indeed, forming $\mathbf{A}^\top[n-1]\mathbf{k}(\mathbf{x}_{\mathcal{S}_{n-1}},\mathbf{x}_n)$ requires Br multiplications for the matrix-vector product, and $\mathcal{O}(BD)$ for the evaluation of B kernels in $\mathbf{k}(\mathbf{x}_{\mathcal{S}_{n-1}},\mathbf{x}_n)$; the matrix-vector product that remains for obtaining $\mathbf{q}[n]$ requires r^2 additional multiplications. Overall, running OK-FEB on N data and with a value of ϵ such that $\tilde{N} \leq N$ data are used for updates requires $\mathcal{O}(\tilde{N}(Br(B+r)+r^3)+N(B(D+r)+r^2))$. Alternatively, tuning ϵ such that $\Pr\{\ell_n > \epsilon\} = \mathbb{E}[\tilde{N}/N] = \gamma$ yields an expected complexity $\mathcal{O}(N(Br(\gamma(B+r)+1) + (\gamma r+1)r^2 + BD))$. As simulation tests will corroborate, the budget parameter B can be chosen as $B = cr$ with $c \in [1.5, 5]$. Thus, we can simplify the derived complexity order as $\mathcal{O}(Nr^2(\gamma r+1) + NDr)$.

6 Stability analysis of kernel approximation

In this section, the effect of low-rank approximation of the lifted vectors on kernel-matrix approximation as well as kernel-based classification and regression is analytically quantified. To this end, recall that given $\{\mathbf{x}_\nu\}_{\nu=1}^N$, the virtual subspace obtained by running OK-FEB is $\bar{\mathbf{L}} = \mathbf{\Phi}_{\mathcal{S}}\mathbf{A} \in \mathbb{R}^{\bar{D}\times r}$, and the corresponding projection coefficients are $\mathbf{Q}_N$.

By defining the random variables $e_i := \|\boldsymbol{\phi}(\mathbf{x}_i) - \hat{\boldsymbol{\phi}}(\mathbf{x}_i)\|^2_{\mathcal{H}} = \|\boldsymbol{\phi}(\mathbf{x}_i) - \bar{\mathbf{L}}\mathbf{q}_i\|^2_{\mathcal{H}}$ capturing the LS error, we have the following result.

Proposition 4. *If the random variables $e_i \in [0\ \ 1]$ are i.i.d. with mean $\bar{e} := \mathbb{E}[e_i]$, then for kernels satisfying $|\kappa(\mathbf{x}_i, \mathbf{x}_j)| \leq 1$, the matrix $\mathbf{K} = \boldsymbol{\Phi}^\top\boldsymbol{\Phi}$ can be approximated by $\hat{\mathbf{K}} := \hat{\boldsymbol{\Phi}}^\top\hat{\boldsymbol{\Phi}}$, and with probability at least $1 - 2e^{-2Nt^2}$, it holds that*

$$\frac{1}{N}\|\mathbf{K} - \hat{\mathbf{K}}\|^2_F \leq \sqrt{\bar{e} + t}\,(\sqrt{\bar{e} + t} + 2)\,. \tag{28}$$

Proof: Upon defining $\bar{\mathbf{E}} := \hat{\boldsymbol{\Phi}} - \boldsymbol{\Phi}$, one can write

$$\begin{aligned}
\|\mathbf{K} - \hat{\mathbf{K}}\|_F &= \|\boldsymbol{\Phi}^\top\boldsymbol{\Phi} - \hat{\boldsymbol{\Phi}}^\top\hat{\boldsymbol{\Phi}}\|_F \\
&= \|\boldsymbol{\Phi}^\top\boldsymbol{\Phi} - (\boldsymbol{\Phi} + \bar{\mathbf{E}})^\top(\boldsymbol{\Phi} + \bar{\mathbf{E}})\|_F \\
&= \|2\bar{\mathbf{E}}^\top\boldsymbol{\Phi} + \bar{\mathbf{E}}^\top\bar{\mathbf{E}}\|_F \\
&\leq 2\|\bar{\mathbf{E}}\|_F\|\boldsymbol{\Phi}\|_F + \|\bar{\mathbf{E}}\|^2_F \qquad (29a) \\
&\leq 2\sqrt{N}\|\bar{\mathbf{E}}\|_F + \|\bar{\mathbf{E}}\|^2_F \qquad (29b)
\end{aligned}$$

where in (29a) we used the triangle inequality for the Frobenious norm along with the property $\|\mathbf{BC}\|_F \leq \|\mathbf{B}\|_F\|\mathbf{C}\|_F$, and (29b) holds because for, e.g., radial kernels satisfying $|\kappa(\mathbf{x}_i, \mathbf{x}_j)| \leq 1$, we have

$$\|\boldsymbol{\Phi}\|_F := \sqrt{\mathrm{tr}(\boldsymbol{\Phi}^\top\boldsymbol{\Phi})} = \sqrt{\sum_{i=1}^{N}\kappa(\mathbf{x}_i, \mathbf{x}_i)} \leq \sqrt{N}\,.$$

Furthermore, since $\|\bar{\mathbf{E}}\|_F := \sqrt{\sum_{i=1}^{N} e_i}$, and $\bar{e}_N := (1/N)\sum_{i=1}^{N} e_i$ with $e_i \in [0, 1]$, Hoeffding's inequality yields

$$\Pr\Big(\bar{e}_N - \bar{e} \geq t\Big) \leq e^{-2Nt^2}$$

which in turn implies

$$\Pr\Big(\frac{1}{N}\|\bar{\mathbf{E}}\|^2_F \geq \bar{e} + t\Big) = \Pr\Big(\bar{e}_N \geq \bar{e} + t\Big) \leq e^{-2Nt^2}\,. \tag{30}$$

Finally, taking into account (29b), it follows that with probability at least $1 - 2e^{-2Nt^2}$, we have

$$\|\mathbf{K} - \hat{\mathbf{K}}\|_F \leq N(2\sqrt{\bar{e} + t} + (\bar{e} + t))\,. \tag{31}$$

■

Proposition 4 essentially bounds the kernel approximation mismatch based on how well the projection onto the subspace approximates the lifted data $\boldsymbol{\phi}(\mathbf{x})$.

Remark 3. Consider now decomposing the kernel matrix as

$$\begin{aligned}
\hat{\mathbf{K}} := \hat{\boldsymbol{\Phi}}^\top\hat{\boldsymbol{\Phi}} &= (\bar{\mathbf{L}}\mathbf{Q})^\top(\bar{\mathbf{L}}\mathbf{Q}) = \mathbf{Q}^\top\mathbf{A}^\top\boldsymbol{\Phi}_{\mathcal{S}}^\top\boldsymbol{\Phi}_{\mathcal{S}}\mathbf{A}\mathbf{Q} \\
&= \mathbf{Q}^\top\mathbf{A}^\top\mathbf{K}_{\mathcal{S}}\mathbf{A}\mathbf{Q} = \mathbf{Z}^\top\mathbf{Z}
\end{aligned} \tag{32}$$

where matrix $\mathbf{Z} := \mathbf{K}_{\mathcal{S}}^{1/2}\mathbf{A}\mathbf{Q}$ has size $|\mathcal{S}| \times N$, and $\mathcal{S}$ denotes the budgeted SV set. This factorization of $\hat{\mathbf{K}}$ could have resulted from a linear kernel over the $|\mathcal{S}| \times 1$ training data vectors forming the N columns of $\mathbf{Z}$. Thus, for kernel-based tasks such as kernel classification, regression, and clustering applied to large datasets, we can simply map the $D \times N$ data $\mathbf{X}$ to the corresponding features $\mathbf{Z}$ trained via the proposed solvers, and then simply rely on *fast linear* learning methods to approximate the original kernel-based learning task; that is to approximate the function $f(\mathbf{x}) = \sum_{i\in\mathcal{S}} c_i\kappa(\mathbf{x}, \mathbf{x}_i)$ by the linear function $g(\mathbf{z}) = \mathbf{w}^\top\mathbf{z}$ expressed via the extracted features. Since linear pattern recognition tasks incur complexity $\mathcal{O}(NB^2)$, they scale extremely well for large datasets (with $N \gg$), compared to kernel SVM that incurs complexity $\mathcal{O}(N^3)$. Furthermore, in the testing phase, evaluation of function $f(\mathbf{x})$ requires $\kappa(\mathbf{x}_\nu, \mathbf{x}_i)$ for $\forall i \in \mathcal{S}$ to be evaluated at complexity $\mathcal{O}(|\mathcal{S}|D)$, where $|\mathcal{S}|$ is the number of SVs that typically grows with N. In contrast, if approximated by the linear $g(\mathbf{z})$, function evaluation requires $\mathcal{O}(BD + Br)$ operations including the feature extraction and function evaluation. Setting the budget B to 1.5 to 5 times the rank parameter r, our complexity is of order $\mathcal{O}(rD + r^2)$, which represents a considerable decrease over $\mathcal{O}(|\mathcal{S}|D)$.

Subsequently, we wish to quantify how the performance of linear classification and regression based on the features $\mathbf{K}_{\mathcal{S}}^{1/2}\mathbf{A}\mathbf{Q}$ compares to the one obtained when training with the exact kernel matrix $\mathbf{K}$.

6.1 Stability analysis for kernel-based classification

Kernel-based classifiers using SVM solve the minimization [29, p. 205]

$$\begin{aligned}
\boldsymbol{\alpha}^* = \arg\min_{\boldsymbol{\alpha}}\ & \frac{1}{2}\boldsymbol{\alpha}^\top\mathbf{Y}\mathbf{K}\mathbf{Y}\boldsymbol{\alpha} - \mathbf{1}^\top\boldsymbol{\alpha} \\
\text{s.t. } & \mathbf{y}^\top\boldsymbol{\alpha} = 0 \\
& \mathbf{0} \leq \boldsymbol{\alpha} \leq \frac{C}{N}\mathbf{1}_N
\end{aligned} \tag{33}$$

where $\mathbf{Y}$ is the diagonal matrix with the i-th label y_i as its i-th diagonal entry, $\mathbf{y}^\top := [y_1, y_2, ..., y_N]$, and $\mathbf{1}_N$ is a $n \times 1$ vector of 1's. The solution (33) corresponds to the dual variables of the primal optimization problem, which yields

$$\bar{\mathbf{w}}^* = \arg\min_{\bar{\mathbf{w}}\in\mathbb{R}^{\bar{D}}} \frac{1}{2}\|\bar{\mathbf{w}}\|^2_{\mathcal{H}} + \frac{C}{N}\sum_{i=1}^{N}\max\{0, 1 - y_i\bar{\mathbf{w}}^\top\boldsymbol{\phi}(\mathbf{x}_i)\}\,. \tag{34}$$

Here, parameter C controls the trade-off between maximization of the margin $1/\|\mathbf{w}\|_{\mathcal{H}}$, and the minimization of the misclassification penalty, while the solution of (34) can be expressed as $\bar{\mathbf{w}}^* = \sum_{i=1}^{N}\alpha_i^* y_i\boldsymbol{\phi}(\mathbf{x}_i)$ [29, p.187].

Exploiting the reduced memory requirement offered through the low-rank approximation of the kernel matrix via OK-FEB, the dual problem can be approximated as

$$\begin{aligned}
\hat{\boldsymbol{\alpha}}^* = \arg\min_{\boldsymbol{\alpha}}\ & \frac{1}{2}\boldsymbol{\alpha}^\top\mathbf{Y}\hat{\mathbf{K}}\mathbf{Y}\boldsymbol{\alpha} - \mathbf{1}^\top\boldsymbol{\alpha} \\
\text{s.t. } & \mathbf{y}^\top\boldsymbol{\alpha} = 0 \\
& \mathbf{0} \leq \boldsymbol{\alpha} \leq \frac{C}{N}\mathbf{1}_N\,.
\end{aligned} \tag{35}$$

Viewing $\hat{\mathbf{K}}$ as a linear kernel matrix over $\{\hat{\boldsymbol{\phi}}(\mathbf{x}_i)\}$s (cf. Remark 3), similar to (33), the minimization (35) can be rewritten in the primal form as

$$\hat{\bar{\mathbf{w}}}^* = \arg\min_{\bar{\mathbf{w}}} \frac{1}{2}\|\bar{\mathbf{w}}\|^2_{\mathcal{H}} + \frac{C}{N}\sum_{i=1}^{N}\max\{0, 1 - y_i\bar{\mathbf{w}}^\top\hat{\boldsymbol{\phi}}(\mathbf{x}_i)\} \tag{36}$$

for which we have $\hat{\bar{\mathbf{w}}}^* = \sum_{i=1}^{N}\hat{\alpha}_i^* y_i\hat{\boldsymbol{\phi}}(\mathbf{x}_i)$. Upon defining the random variable $\xi_i := \|\boldsymbol{\phi}(\mathbf{x}_i) - \hat{\boldsymbol{\phi}}(\mathbf{x}_i)\|_{\mathcal{H}}$ with expected value $\bar{\xi} := \mathbb{E}[\xi_i]$, the following proposition quantifies the gap between $\bar{\mathbf{w}}^*$ and $\hat{\bar{\mathbf{w}}}^*$.

Proposition 5. *If variables $\xi_i \in [0 \;\; 1]$ are i.i.d., the mismatch between the linear classifiers given by* (34) *and* (36) *can be bounded, and with probability at least $1-e^{-2Nt^2}$, we have*

$$\|\Delta\mathbf{w}\|_{\mathcal{H}}^2 := \|\bar{\mathbf{w}}^* - \hat{\bar{\mathbf{w}}}^*\|_{\mathcal{H}}^2 \leq 2C^{3/2}\Big(\bar{\xi}+t\Big) \; . \tag{37}$$

Proof: It clearly holds that

$$\begin{aligned}\|\Delta\mathbf{w}\|_{\mathcal{H}}^2 &\leq \frac{C}{N}(\|\bar{\mathbf{w}}^*\| + \|\hat{\bar{\mathbf{w}}}\|)\sum_{i=1}^{N}\|\boldsymbol{\phi}(\mathbf{x}_i) - \hat{\boldsymbol{\phi}}(\mathbf{x}_i)\|_{\mathcal{H}} \\ &\leq \frac{2C^{3/2}}{N}\sum_{i=1}^{N}\|\boldsymbol{\phi}(\mathbf{x}_i) - \hat{\boldsymbol{\phi}}(\mathbf{x}_i)\|_{\mathcal{H}} \\ &\leq 2C^{3/2}(\bar{\xi}+t) \end{aligned} \tag{38}$$

where the first inequality relies on the strong convexity of (34), (36), and the fact that $\|\bar{\mathbf{w}}^*\|_{\mathcal{H}} \leq \sqrt{C}$ and $\|\hat{\bar{\mathbf{w}}}^*\|_{\mathcal{H}} \leq \sqrt{C}$ [32]; while the second inequality holds with probability at least $1-e^{-2Nt^2}$ using Hoeffding's inequality for $\bar{\xi}_N := (1/N)\sum_{i=1}^{N}\xi_i$. ■

Note that under the i.i.d. assumption on $e_i := \|\boldsymbol{\phi}(\mathbf{x}_i) - \hat{\boldsymbol{\phi}}(\mathbf{x}_i)\|_{\mathcal{H}}^2$, random variables ξ_i are also i.i.d., rendering the conditions of Propositions 4 and 5 equivalent.

Next, we study the performance of linear SVM trained on the set $\{\mathbf{z}_i, y_i\}_{i=1}^N$, where $\mathbf{z}_i := \mathbf{K}_{\mathcal{S}}^{1/2}\mathbf{A}\mathbf{q}_i$; that is, the linear function $g(\mathbf{z}) = \mathbf{w}^\top\mathbf{z}$ is learned by finding

$$\mathbf{w}^* = \arg\min_{\mathbf{w}\in\mathbb{R}^r} \frac{1}{2}\|\mathbf{w}\|^2 + \frac{C}{N}\sum_{i=1}^{N}\max\{0, 1-y_i\mathbf{w}^\top\mathbf{z}_i\} \; . \tag{39}$$

The following result asserts that the classifiers learned through (39) and (36) can afford identical generalization capabilities.

Proposition 6. *The generalization capability of classifiers* (36) *and* (39) *is identical, in the sense that $\hat{\bar{\mathbf{w}}}^{*\top}\hat{\boldsymbol{\phi}}(\mathbf{x}) = \mathbf{w}^{*\top}\mathbf{z}$.*

Proof: Since for the low-rank approximation of the kernel matrix we have $\hat{\mathbf{K}} = \mathbf{Z}^\top\mathbf{Z}$, then (35) and (39) are equivalent, and consequently $\mathbf{w}^* = \sum_{i=1}^{N}\hat{\alpha}_i^* y_i\mathbf{z}_i$. Now, one can further expand $\hat{\bar{\mathbf{w}}}^{*\top}\hat{\boldsymbol{\phi}}(\mathbf{x})$ and $\mathbf{w}^{*\top}\mathbf{z}$ to obtain

$$\hat{\bar{\mathbf{w}}}^{*\top}\hat{\boldsymbol{\phi}}(\mathbf{x}) = \sum_{i=1}^{N}\hat{\alpha}_i^* y_i\hat{\boldsymbol{\phi}}(\mathbf{x}_i)^\top\hat{\boldsymbol{\phi}}(\mathbf{x}) = \sum_{i=1}^{N}\hat{\alpha}_i^* y_i\mathbf{q}_i\mathbf{A}\hat{\boldsymbol{\Phi}}_{\mathcal{S}}^\top\hat{\boldsymbol{\Phi}}_{\mathcal{S}}\mathbf{A}\mathbf{q}$$

and

$$\mathbf{w}^{*\top}\mathbf{z} = \sum_{i=1}^{N}\hat{\alpha}_i^* y_i\mathbf{z}_i^\top\mathbf{z} = \sum_{i=1}^{N}\hat{\alpha}_i^* y_i\mathbf{q}_i\mathbf{A}\hat{\boldsymbol{\Phi}}_{\mathcal{S}}^\top\hat{\boldsymbol{\Phi}}_{\mathcal{S}}\mathbf{A}\mathbf{q}$$

where the equivalence follows readily. ■

In addition to markedly reduced computational cost when utilizing linear (L)SVM, our novel classifier can also be efficiently trained online [32] as new data becomes available (or iteratively when the entire datasets can not be stored in memory which necessitates one-by-one acquisition). In this case, the proposed OK-FEB in Algorithm 4 can be run in parallel with the online classifier training, an attribute most suitable for big data applications.

6.2 Stability analysis for kernel-based regression

Consider the kernel-based ridge regression task on the dataset $\{\mathbf{x}_i, y_i\}_{i=1}^N$, namely

$$\min_{\boldsymbol{\beta}} \frac{1}{N}\|\mathbf{y} - \mathbf{K}\boldsymbol{\beta}\|_2^2 + \lambda\boldsymbol{\beta}^\top\mathbf{K}\boldsymbol{\beta} \tag{40}$$

which admits the closed-form solution $\boldsymbol{\beta}^* = (\mathbf{K}+\lambda N\mathbf{I})^{-1}\mathbf{y}$ [29, p.251]. Alleviating the memory requirement of order $\mathcal{O}(N^2)$ through low-rank approximation of the kernel matrix, the kernel-based ridge regression can be approximated by

$$\min_{\boldsymbol{\beta}} \frac{1}{N}\|\mathbf{y} - \hat{\mathbf{K}}\boldsymbol{\beta}\|_2^2 + \lambda\boldsymbol{\beta}^\top\hat{\mathbf{K}}\boldsymbol{\beta} \tag{41}$$

whose solution is given as $\hat{\boldsymbol{\beta}}^* = (\hat{\mathbf{K}}+\lambda N\mathbf{I})^{-1}\mathbf{y}$. The following proposition bounds the mismatch between $\boldsymbol{\beta}^*$ and $\hat{\boldsymbol{\beta}}^*$.

Proposition 7. *If the random variables $e_i \in [0 \;\; , \;\; 1]$ are i.i.d., and $|y_i| \leq B_y$ for $i = 1, 2, ..., N$, with probability at least $1-2e^{-2Nt^2}$ we have*

$$\|\boldsymbol{\beta}^* - \hat{\boldsymbol{\beta}}^*\|^2 \leq \frac{M}{\lambda^2}\sqrt{\bar{e}+t}(\sqrt{\bar{e}+t}+2) \; . \tag{42}$$

Proof: Following [7], we can write

$$\begin{aligned}\boldsymbol{\beta}^* - \hat{\boldsymbol{\beta}}^* &= (\mathbf{K}+\lambda N\mathbf{I})^{-1}\mathbf{y} - (\hat{\mathbf{K}}+\lambda N\mathbf{I})^{-1}\mathbf{y} \\ &= -\Big((\hat{\mathbf{K}}+\lambda N\mathbf{I})^{-1}(\mathbf{K}-\hat{\mathbf{K}})(\mathbf{K}+\lambda N\mathbf{I})^{-1}\Big)\mathbf{y}\end{aligned}$$

where we have used the identity $\hat{\mathbf{P}}^{-1} - \mathbf{P}^{-1} = -\mathbf{P}^{-1}(\hat{\mathbf{P}} - \mathbf{P})\hat{\mathbf{P}}^{-1}$, which holds for any invertible matrices $\mathbf{P}$ and $\hat{\mathbf{P}}$. Taking the ℓ_2-norm of both sides and using the Cauchy-Schwarz inequality, we arrive at

$$\begin{aligned}\|\boldsymbol{\beta}^* - \hat{\boldsymbol{\beta}}^*\| &\leq \|(\mathbf{K}+\lambda N\mathbf{I})^{-1}\|\|\mathbf{K}-\hat{\mathbf{K}}\|\|(\hat{\mathbf{K}}+\lambda N\mathbf{I})^{-1}\|\|\mathbf{y}\| \\ &\leq \frac{\|\mathbf{K}-\hat{\mathbf{K}}\|NB_y}{\lambda_{\min}(\mathbf{K}+\lambda N\mathbf{I})\lambda_{\min}(\hat{\mathbf{K}}+\lambda N\mathbf{I})} \\ &\leq \frac{B_y\|\mathbf{K}-\hat{\mathbf{K}}\|_2}{\lambda^2 N} \; .\end{aligned} \tag{43}$$

Using the inequality $\|\mathbf{P}\|_2 \leq \|\mathbf{P}\|_F$ along with Proposition 4, yields the bound with probability $1-2e^{-2Nt^2}$. ■

7 Numerical tests

Performance of our proposed algorithm is assessed in this section using synthetic and real datasets on kernel-matrix approximation as well as classification and regression tasks. The synthetic test involves generating two equiprobable classes of 3×1 data vectors $\{\mathbf{x}_\nu\}$, with each data vector uniformly drawn from the surface of a sphere centered at the origin with radius $R_{c1} = 1$ or $R_{c2} = 2$, depending on whether its label y_ν equals 1 or -1 respectively. A noise component drawn from the Gaussian distribution $\mathcal{N}(\mathbf{0}_{D\times 1}, \sigma^2\mathbf{I}_{D\times D})$ is added to each $\mathbf{x}_\nu$, with the variance σ^2 controlling the overlap between the two classes. Linear classifiers can not correctly classify data generated in this manner; hence, a kernel-based classifier is well motivated. This dataset is used to validate performance of the proposed OK-FEB in kernel-matrix approximation and generalization capability of the trained linear classifier. The size and specifications of the dataset used are listed in Table 1. The datasets

are accessible in the LIBSVM website[1] or the UCI machine learning repository [2].

TABLE 1
Specifications of datasets.

dataset	task	D	N
synthetic	classification	3	70K
IJCNN	classification	22	140K
Adult	classification	123	32K
Cifar	classification	3072	50K
MNIST	classification	786	60K
COD-RNA	classification	8	59K
CADATA	regression	8	20.6K
Slice	regression	384	53.5K
Year	regression	90	463.7K

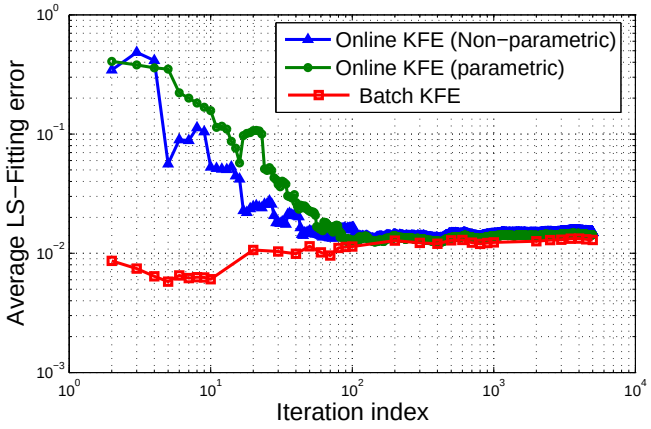

Fig. 1. Average LS-fit versus iteration index for the synthetic dataset.

7.1 Kernel-based feature extraction: Batch versus online

Performance of the proposed Algorithms 1, 2 and 3 on solving the minimization (10) is compared empirically. The test is carried out for the synthetically generated data arriving in streaming mode with $\nu = 1, 2, \ldots, 5,000$. The online schemes can solve the problem on-the-fly, while the batch Algorithm 1 is also employed to solve (10) for all n. We compare the overall LS fit given by the subspace update $\bar{\mathbf{L}}[n]$ using the three different solvers across time (iteration) index n. Gaussian kernel $\kappa(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\|\mathbf{x}_i - \mathbf{x}_j\|_2^2/\gamma)$ is used with $\gamma = 100$ set to the standard deviation of average pairwise squared distances of the data vectors. The parameters for the OK-FE solvers are chosen as $\mu_{n,L} = 10/n$, $\mu_{n,A} = 1000/n^2$, $\lambda = 10^{-3}$, and the maximum number of iterations in the batch solver is set to $K_{\max} = 50$.

Figure 1 depicts how stochastic low-complexity updates of $\mathbf{A}$ in the online solvers ensure convergence of the average LS cost to the high-complexity batch solution. In addition, Fig. 2 plots the evolution of the average LS fit cost across iterations for different choices of parameters (r, B) in the proposed OK-FEB solver.

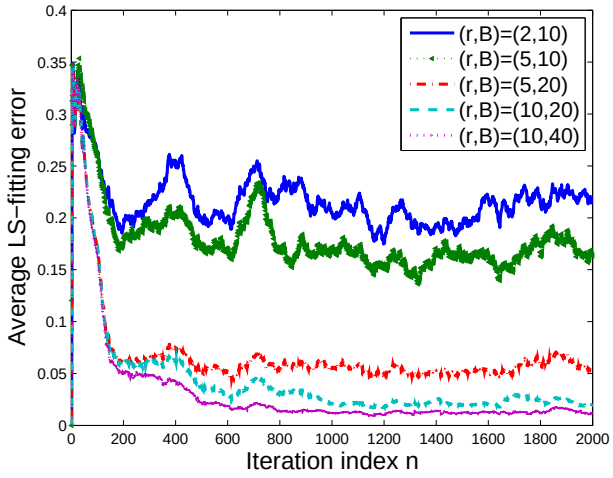

Fig. 2. Average LS-fit for different choices of (r, B) versus iteration index for the synthetic dataset.

7.2 Kernel-matrix approximation

To quantify the kernel matrix approximation accuracy, $N^{-1}\|\mathbf{K} - \hat{\mathbf{K}}\|_F$ provided by the novel approach is measured here on real datasets, and is compared with that of kernel-matrix approximation using four different schemes: i) plain Nystrom approximation (random selection of "landscape" data points) [38]; ii) Nystrom approximation with nonuniform (probabilistic) data point selection [11]; iii) Improved Nystrom (where landscape points are chosen via k-means) [40]; and, iv) rank-r approximation of $\mathbf{K}$ via singular-value-decomposition (SVD) as the benchmark. Since SVD incurs complexity $\mathcal{O}(N^3)$, it is computationally prohibitive for such large data sets. For this reason, we have approximated $N^{-1}\|\mathbf{K} - \hat{\mathbf{K}}_{SVD}\|_F$ of the full $N \times N$ kernel matrix by that on a submatrix of size $10^4 \times 10^4$, averaged over several selections. While the proposed OK-FEB only requires acquiring one data vector at a time, the other comparable methods are developed for batch settings, where the entire dataset must be loaded in memory, a requirement not always met in practice. However, to meet the memory requirement of such algorithms, the experiments are performed here on a 16-core node machine with 2.50 GHz CPU and 225GB of RAM.

The budgeted subspace learning in the OK-FEB algorithm is highly efficient, and thanks to its fast convergence, only one pass over a subset of training data vectors is sufficient for stationary or batch datasets. Thus, subspace learning is carried over $2-5\%$ of the set, while the rest of the features are simply acquired by projecting onto the identified subspace. This suggests that in stationary batch settings, the complexity of the proposed algorithm is comparable to that of the Nystrom approximation, which is minimal in such scenaria, while exhibiting considerable improvement in performance.

1. http://www.csie.ntu.edu.tw/~cjlin/libsvmtools/
2. http://www.ics.uci.edu/~mlearn/

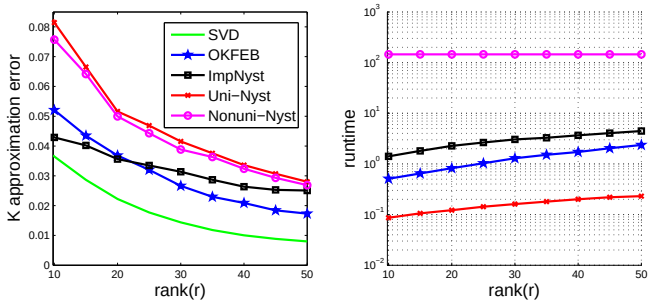

Fig. 3. Normalized kernel matrix approximation error, and run time for the Adult dataset.

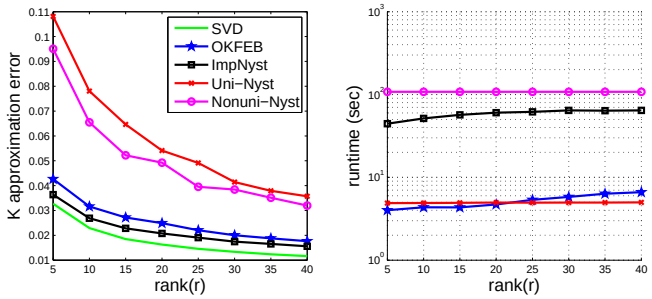

Fig. 4. Normalized kernel matrix approximation error, and run time for the CIFAR dataset.

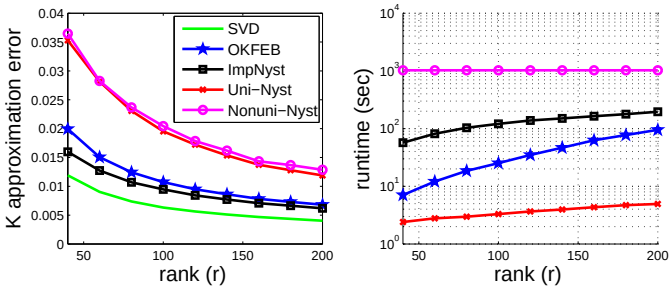

Fig. 5. Normalized kernel matrix approximation error, and run time for the MNIST dataset $(B = 1.2r)$.

TABLE 2
Parameters for the classification and regression tests

Dataset	r	B	γ	ϵ	C
synthetic	15	1.5r	6	0.5	1
Adult	30	$2r$	20	0.5	100
IJCNN	200	1.5r	1	0.1	1

7.3 OK-FEB for classification and regression

In this subsection, the generalization capability of the batch linear classification and regression tasks based on the features $\mathbf{Z}$ returned by OK-FEB are compared with that of their batch kernel-based SVM counterparts. Training parameters are set through cross validation, and are reported in Table II. The experiments of this section are performed on a machine of 2.50 GHz quad-core CPU, and 16GB of RAM. In Table III, the out-of-sample classification accuracy of the kernel and linear tasks as well as the training and testing times are reported.

TABLE 3
Run time for "OK-FEB+LSVM" versus "K-SVM"

Dataset	Adult	IJCNN	Synthetic
$t_{\text{OK-FEB}}$	5.43	51.06	14.46
$t_{\text{LSVM-train}}$(s)	0.36	15.45	16.72
$t_{\text{total-train}}$(s)	**5.79**	**66.51**	**31.18**
$t_{\text{KSVM-train}}$(s)	34.20	113.82	62.31
$t_{\text{LSVM-test}}$(s)	**1.12**	**3.27**	3.31
$t_{\text{KSVM-test}}$(s)	4.8	10.07	1e-5
$\text{acc}_{\text{OK-FEB+LSVM}}$	84.01%	96.70	73.45 %
acc_{KSVM}	85.15%	98.89	73.49

TABLE 4
Parameters for online classification and regression tests

Dataset	r	B	γ	C
Adult	50	$1.2r$	20	10
COD-RNA	20	$1.5r$	600	10
CADATA	5	$1,5r$	7×10^7	0.01
Slice	10	$1.2r$	50	0.01
Year	10	$1.2r$	5×10^7	0.01

7.4 Online regression and classification

In this section, we compare the performance of competing online solvers including (unbudgeted) Perceptron [14], (unbudgeted) Norma [18], (unbudgeted) online gradient descent (OGD) [32], and budgeted online gradient descent (BOGD) [35], Forgetron [9], Projectron [26], and budgeted Passive-Aggressive algorithm (BPA) [37] with the proposed OK-FEB, where the acquired features $\mathbf{z}_n$ are fed to online linear Pegasus [32] and regularized-LMS solver for the classification and regression tasks, respectively. The parameter values used for each dataset are reported in Table IV, and are set via cross validation. The censoring threshold ϵ is set using a moving-average of LS-error values for the past 100 data vectors.

Classification and regression accuracy as well as run time of the algorithms are plotted versus iteration index. Perceptron, Norma, and OGD are unbudgeted algorithms, and the SV set grows as iteration index increases in Figs. 6-10. Although the accuracy of these algorithms can serve as a benchmark, their run time grows the fastest. Thus, for the "Year" dataset (very large N), the mentioned algorithms are run only over 10% of the data vectors. As these tests demonstrate, among the budgeted algorithms, OK-FEB reliably approximates the kernel function through the extracted features, and thus offers a more accurate classification or regression approximation when compared to existing alternatives.

8 Concluding Remarks

Low-complexity feature extraction algorithms were introduced in this paper to markedly improve the performance of kernel-based learning methods applied to large-scale datasets. Our approach begins with a generative model where data are mapped to a high (possibly infinite) dimensional space, where they can be well approximated

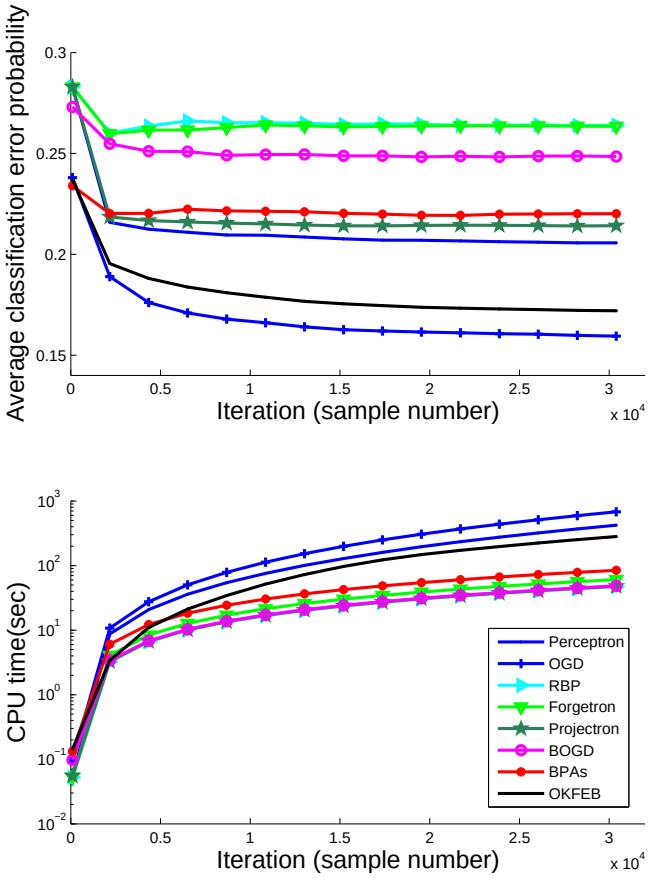

Fig. 6. Classification test using the 'Adult' dataset.

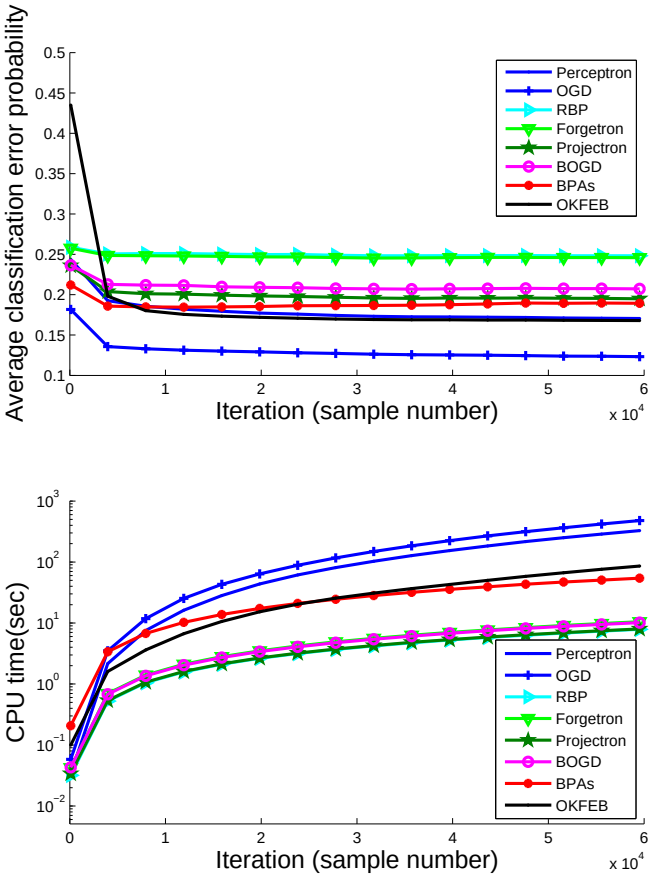

Fig. 7. Classification test using the 'COD-RNA' dataset.

to lie close to a *linear* low-rank subspace, the tracking of which enables effective feature extraction on a budget. The extracted features can be used by fast linear classification or regression solvers to efficiently perform the learning tasks.

Offline and online methods for solving the subspace learning task were developed, and their convergence was studied analytically. To keep the complexity and memory requirements within affordable levels, *budgeted* versions of the proposed algorithms were devised, in which the number of stored data vectors is restricted to a prescribed budget. Further analysis provided performance bounds on the quality of the kernel matrix approximation, as well as the precision with which kernel-based classification and regression tasks can be approximated by leveraging budgeted online subspace-learning and feature-extraction tasks. Finally, classification and regression tests on synthetic and real datasets demonstrated the efficiency of the proposed OK-FEB method with respect to competing alternatives, in terms of both accuracy and run time.

REFERENCES

[1] F. Bach, "Sharp analysis of low-rank kernel matrix approximations," Technical report. Preprint: arXiv:1208.2015, 2012.
[2] D. P. Bertsekas, *Nonlinear Programming*, Athena Scientific, 1999.
[3] D. Berberidis, V. Kekatos, and G.B. Giannakis, "Online Censoring for Large-Scale Regressions with Application to Streaming Big Data," *arXiv:1507.07536v1*, 2015 (submitted)
[4] A. Bordes, S. Ertekin, J. Weston, and L. Bottou, "Fast kernel classifiers with online and active learning," *J. of Mach. Learn. Res.*, vol. 6, pp. 1579-1619, 2005.
[5] G. Cavallanti, N. Cesa-Bianchi, and C. Gentile, "Tracking the best hyperplane with a simple budget perceptron," *Machine Learning*, vol. 69, pp. 143-167, 2007.
[6] C. C. Chang and C. J. Lin, "LIBSVM: A library for support vector machines," *ACM Trans. Intelligent Systems and Technology*, vol. 2, no. 3, pp. 27:1-27:27, 2011.
[7] C. Cortes, M. Mohri, and A. Talwalkar, "On the impact of kernel approximation on learning accuracy," *Proc. of Intl. Conf. on Artif. Intel. and Stat.*, pp. 113-120, Sardinia, Italy, May 2010.
[8] B. Dai, B. Xie, N. He, Y. Liang, A. Raj, M. F. Balcan, and L. Song, "Scalable kernel methods via doubly stochastic gradients," *Proc. of Advances in Neural Inf. Proc. Systems*, pp. 3041-3049, Montreal, Canada, Dec. 2014.
[9] O. Dekel, S. Shalev-Shwartz, and Y. Singer, "The forgetron: A kernel-based perceptron on a budget," *SIAM J. on Computing*, vol. 37, no. 5, pp. 1342-1372, 2008.
[10] O. Dekel, and Y. Singer, "Support vector machines on a budget," *Proc. of Advances in Neural Inf. Proc. Systems*, pp. 345-352, Vancouver, Canada, Dec. 2007.
[11] P. Drineas, M. W. Mahoney, "On the Nystrom method for approximating a Gram matrix for improved kernel-based learning," *J. of Mach. Learn. Res.*, vol. 6, pp. 2153-2175, Dec. 2005.
[12] R. E. Fan, K. W. Chang, C. J. Hsieh, X. R. Wang, and C. J. Lin, "LIBLINEAR: A Library for Large Linear Classification," *J. of Mach. Learn. Res.*, vol. 9, pp. 1871-1874, 2008.
[13] S. Fine, and K. Scheinberg, "Efficient SVM training using low-rank kernel representations," *J. of Mach. Learn. Res.*, vol. 2, pp. 243-264, 2002.
[14] Y. Freund, R. E. Schapire, "Large margin classification using the

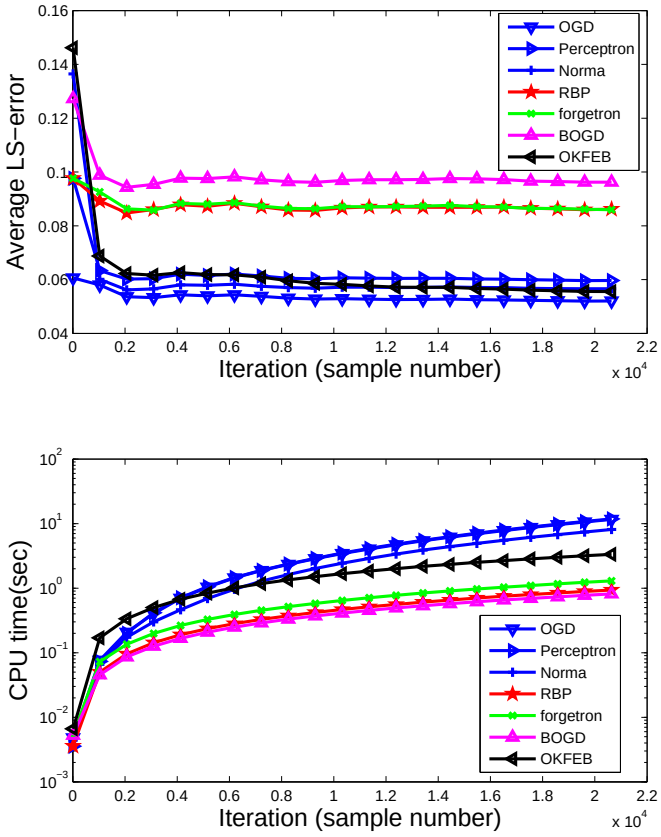

Fig. 8. Regression test using the 'CADATA' dataset.

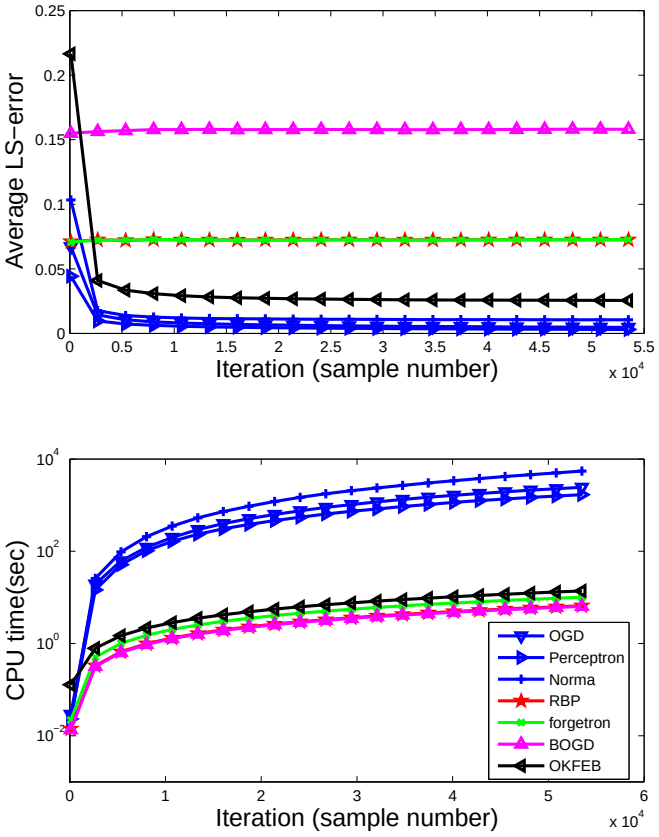

Fig. 9. Regression test using the 'Slice' dataset.

perceptron algorithm," *Machine Learning*, vol. 37, No. 3, pp. 277-296, 1999.

[15] Y. Hastie, R. Tibshirani, and J. Friedman, *The Elements of Statistical Learning*, Springer, 2008.

[16] P. Honeine, "Online kernel principal component analysis: A reduced-order model," *IEEE Trans. on Patt. Anal. and Mach. Intel.*, vol. 34, no. 9 pp. 1814-1826, 2012.

[17] K. Kim, M. Franz, and B. Scholkopf, "Iterative principal component analysis for image modeling," *IEEE Trans. on Patt. Anal. and Mach. Intl.*, vol. 27, no. 9, pp. 1351-1366, 2005.

[18] J. Kivinen, A. Smola, and R. Williamson, "Online learning with kernels," *IEEE Trans. on Sign. Proc.*, vol. 52, no. 8, pp. 2165-2176, August 2004.

[19] S. Kumar, M. Mohri, and A. Talwalkar, "Ensemble Nystrom method," *Proc. of Advances in Neural Inf. Proc. Systems*, pp. 1060-1068, Vancouver, Canada, 2009.

[20] Y. J. Lee, and O.L. Mangasarian, "RSVM: Reduced Support Vector Machines," *Proc. of SIAM Intl. Conf. on Data Mining*, vol. 1, pp. 325-361, Chicago, USA, April 2001.

[21] D. Lopez-Paz, S. Sra, A. Smola, Z. Ghahramani, and B. Scholkopf, "Randomized nonlinear component analysis," *Proc. of Intl. Conf. on Mach. Learn.*, pp. 1359-1367, Beijing, China, June 2014.

[22] J. Lu, S. C. H. Hoi, J. Wang, P. Zhao, Z. Liu, "Large scale online kernel learning," *J. of Mach. Learn. Res.*, 2015, (to appear).

[23] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *J. of Mach. Learn. Res.*, vol. 11, pp. 19-60, March 2011.

[24] M. Mardani, G. Mateos, and G. B. Giannakis, "Dynamic anomalography: Tracking network anomalies via sparsity and low-rank," *IEEE J. Sel. Topics in Sig. Proc.*, vol. 7, no. 1, pp. 50-66, Feb. 2013.

[25] A. Nemirovski, J. Anatoli, L. G. Lan, and A. Shapiro, "Robust stochastic approximation approach to stochastic programming," *SIAM J. on Optim.*, no. 4, vol. 19, pp. 1574-1609, 2009.

[26] F. Orabona, J. Keshet, B. Caputo, "The projectron: A bounded kernel-based perceptron," *Proc. of Intl. Conf. on Mach. Learn.*, pp. 720-727, Chicago, USA, July 2011.

[27] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," *Proc. of Advances in Neural Inf. Proc. Systems*, pp. 1177-1184, Lake Tahoe, USA, Dec. 2013.

[28] B. Recht, M. Fazel, and P. A. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, pp. 471-501, 2010.

[29] B. Scholkopf, and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and beyond*, MIT Press, 2001.

[30] B. Scholkopf, A. Smola, and K. R. Muller, "Nonlinear component analysis as a kernel eigenvalue problem," *Neural Computation*, vol. 10.5, pp. 1299-1319, 1998.

[31] N. Schraudolph, S. Gunter, and S. V. N. Vishwanathan, "Fast iterative kernel principal component analysis," *J. of Mach. Learn. Res.*, vol. 8, pp. 1893-1918, 2008.

[32] S. Shalev-Shwartz, Y. Singer, N. Srebro, and A. Cotter, "Pegasos: Primal estimated sub-gradient solver for SVM," *Mathematical Programming*, Springer, vol. 127, no. 1, pp. 3-30, 2011.

[33] I. W. Tsang, J. T. Kwok, and P. M. Cheung. "Core vector machines: Fast SVM training on very large data sets," *J. of Mach. Learn. Res.*, pp. 363-392, 2005.

[34] S. Van Vaerenbergh, I. Santamaria, W. Liu, and J. C. Principe, "Fixed-budget kernel recursive least-squares," *Proc of Intl. Conf. on Acoustics Speech and Sign. Proc.*, pp. 1882-1885, Dallas, TX, March 2010.

[35] Z. Wang, K. Crammer, and S. Vucetic, "Breaking the curse of kernelization: Budgeted stochastic gradient descent for large scale SVM training," *J. of Mach. Learn. Res.*, vol. 13, pp. 3103-3131, 2012.

[36] Z. Wang, and S. Vucetic, "Online training on a budget of support vector machines using twin prototypes," *Statistical Analysis and*

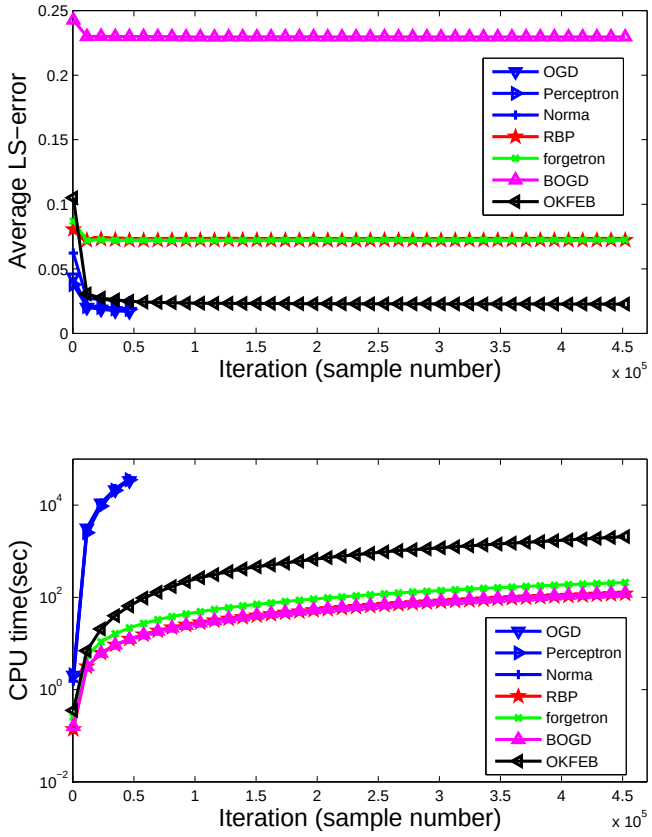

Fig. 10. Regression test using the 'Year' dataset.

Data Mining: The ASA Data Science Journal, vol. 3, no. 3, pp. 149-169, 2010.

[37] Z. Wang, and S. Vucetic, "Online passive-aggressive algorithms on a budget," *J. of Mach. Learn. Res.*, vol. 9, pp. 908-915, 2010.

[38] C. Williams, and M. Seeger, "Using the Nystrom method to speed up kernel machines," *Proc. of Advances on Neural Inf. Proc. Systems*, pp. 682-688, Vancouver, Canada, 2001.

[39] K. Zhang, L. Lan, Z. Wang, and F. Moerchen, "Scaling up kernel SVM on limited resources: A low-rank linearization approach," *Proc. of Intl. Conf. on Artif. Intel. and Stat.*, pp. 1425-1434, La Palma, Canary Islands, April 2012.

[40] K. Zhang, I. Tsang, and J. Kwok, "Improved Nystrom low-rank approximation and error analysis," *Proc. of Intl. Conf. on Mach. Learn.*, pp. 1232-1239, Helsinki, Finland, July 2008.

[41] T. Yang, Y. Li, M. Mahdavi, R. Jin, and Z. Zhou, "Nystrom method vs random Fourier features: A theoretical and empirical comparison," *Proc. of Advances in Neural Inf. Proc. Systems*, pp. 476-484., Lake Tahoe, USA, Dec. 2012.